ADAPTIVE SPHERICAL WAVELETS FOR THE ANGULAR DISCRETISATION OF THE BOLTZMANN TRANSPORT EQUATION

by

Andrew Buchan, Msci

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy and the Diploma of Imperial College

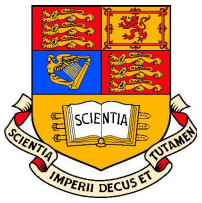

Computational Physics and Geophysics Group
Department of Earth Sciences and Engineering
Imperial College of Science, Technology and Medicine
University of London

October, 2006

To my mother and father, Christine and George for their love and support

Abstract

Three new spherical wavelet bases have been developed to describe the direction of particle travel for neutral particle transport. Two methods are low order linear and quadratic wavelet bases developed on the ideas of Swelden's second generation wavelets. The third is a spectral wavelet basis based on a Chebyshev collocation point method. The wavelet bases are employed to discretise the angular variable of the one speed, linear Boltzmann transport equation (BTE), which is used to describe the transport of neutral particles. The methods are combined with a Streamline Up-wind Petrov Galerkin (SUPG) finite element method which is employed to discretise the transport equation's spatial dimensions.

A comprehensive demonstration of the new angular discretisation methods is given by solving a number of steady state multi-dimensional neutral particle transport problems. These demonstrate the capabilities of the wavelet methods by showing them to produce very accurate solutions in comparison to the established discretisation methods S_N and P_N. The numerical examples also demonstrate the wavelet's ability to accurately approximate difficult flux distributions and to reduce ray-effect distortions forming in their scalar flux solutions (in comparison to S_N).

Also presented in this thesis is a wavelet method for anisotropic adaptive resolution on the unit sphere for approximating the angular variable of the BTE. The method is designed to locate the areas of angle containing high errors and to react correctly by increasing resolution only to these regions on the unit sphere. The self-adaptive method's capabilities are demonstrated by solving a number of steady-state particle transport problems. In particular, problems involving difficult, highly directional fluxes are used to illustrate the method working correctly by its application of high resolution only to the important areas of angle. The numerical examples demonstrate the adaptive wavelets to be highly efficient in approximating the angular flux and to provide very accurate solutions.

A new Krylov-based matrix solver is also developed. This solver uses a new variation of the FGMRES algorithm which is preconditioned with a new hierarchical angular multigrid (MG) method. A demonstration of the preconditioned solver's efficiency using various MG cycles is presented.

Acknowledgements

I first wish to thank my supervisors Professor Christopher Pain, Dr Richard Smedley-Stephenson and Professor Tony Goddard for their support and guidance throughout the duration of this PhD. I have benefited greatly from their expertise and knowledge of their areas in numerical transport theory. I would also like to show my appreciation to Professor Cassiano de Oliveria who helped organise the funding for this PhD and help guide me in the early stages of its research. My thanks also goes out to Dr Matthew Eaton for providing the code RADIANT that has made this work possible. I also wish to thank Mr Adrian Umpleby for his help with the linear solver technologies and Dr Kemal Zilver for his advice on general neutron transport calculations.

I would also like to thank Professor Ron Ackroyd for his encouragement and advice and for showing a particular interest in this work. I also wish to thank Professor Mike Williams for his help and for presenting a series of helpful neutron transport lectures in the early stages of this PhD. My thanks also goes to Dr Stephan Matthai for his help and useful advice.

I also wish to thank my family, in particular my mother and father Christine and George for all their support and help during the difficult time of writing this PhD.

For the funding of this PhD I would like to thank EPSRC and AWE.

Contents

List of Figures

List of Tables

Nomenclature

Latin Characters

Angular expansion index	N
Spherical harmonic angular discretisation	P_N
Simplified spherical harmonic angular discretisation	SP_N
Simplified solid harmonic angular discretisation	SHP_N
Discrete ordinate angular discretisation	S_N
Linear octahedral wavelet angular discretisation	LW_N
Quadratic octahedral wavelet angular discretisation	QW_N
Linear hexahedral wavelet angular discretisation	HLW_N
Quadratic hexahedral wavelet angular discretisation	HQW_N
Chebyshev wavelet angular discretisation	SHW_N
Angular mass matrix	A
Angular Jacobian for x component	A_x
Angular Jacobian for y component	A_y
Angular Jacobian for z component	A_z
Vector of angular Jacobians	$\boldsymbol{A}$
Dot product of angular Jacobians and outward direction to boundary n	A_s
Matrix of right eigenvectors of A_s	R_s
Diagonal matrix of eigenvalues of A_s	Λ_s
Stabilisation matrix	P
Neutron velocity	v
Spatial variable	r
Energy variable	E
Time variable	t
Normal vector to boundary	n
External neutron source	q_{ex}
Scattering neutron source	q_s
Fission neutron source	q_f

Spherical harmonic	$Y_{l,m}$
Conjugate spherical harmonic	$Y^*_{l,m}$
Real part spherical harmonic	Y^e
Imaginary part spherical harmonic	Y^o
Real cosine spherical harmonic	$Y^c_{l,m}$
Real sine spherical harmonic	$Y^s_{l,m}$
Legendre polynomial	P_l
Associated Legendre polynomial	P^m_l
Walsh function	Wal
Chebyshev polynomials	T_k
Chebyshev Lagrangians	h^i_j
Control volume function	M_i
Diagonal matrix of control volume function	$\boldsymbol{M}_i$
Finite element function	N_i
Diagonal matrix of finite element functions	$\boldsymbol{N}_i$
Surface of unit sphere	S^2
Surface of octahedron	O
Subdivided surface of octahedron	O_i
Surface of hexahedron	H
Surface of subdivided hexahedron	H_i
Index set of wavelet collocation points	S_j
Index set of wavelet collocation points	$K(j)$
Index set of wavelet collocation points	$M(j)$
Lebesgue functional set	$L_2(.)$
Spaces of MRA	V_j
Spaces spanned by wavelets	W_j
Wavelet lifting coefficients	$S_{j,m,k}$
Local 2-D plane	P
Krylov subspace	K_i

Search direction space	K_ζ
Preconditioning Matrix	M
Linear system	A
Linear system formed by a discretisation using element of size h_i	A_{h_i}
Right hand side to a linear system	b
Right hand side to a linear system A_{h_i}	b_{h_i}
Approximation to linear system A_{h_i}	x_j
Residual of approximation $b - A_{h_i} x_j$	r_j

Calligraphic Characters

Angular expansion size	$\mathcal{M}$
Spatial discretisation size	$\mathcal{N}$
Angular basis function	$\mathcal{G}_j$
Vector of angular discretised source	$\mathcal{S}$
Matrix of angular discretised scattering/removal term	$\mathcal{H}$
Matrix of space/Angle discretised scattering/removal term	H

Greek Characters

co-latitude angular variable	μ
azimuth angular variable	ω
Angular variable	$\hat{\Omega}$
Reflective angle to $\hat{\Omega}$	$\hat{\Omega}^*$
Angular component in x direction	$\hat{\Omega}_x$
Angular component in y direction	$\hat{\Omega}_y$
Angular component in z direction	$\hat{\Omega}_z$
Scattering angle	μ_0
Coefficient of Legendre expansion of scattering term	σ_{sl}
Scattering variables	Φ_l^m
Scattering variables	$\alpha_j^{e,l,m}$
Scattering variables	$\alpha_j^{o,l,m}$

Total cross-section	σ_t
Capture cross-section	σ_a
Scattering cross-section	σ_s
Fission cross-section	σ_f
Fission energy probability function	χ
Average particle production function	ν
Angular flux	ψ
Scalar flux	ϕ
Radiative flux	$\underline{\phi}$
Pressure flux	$\hat{\phi}$
Even parity variable	ψ^+
Odd parity variable	ψ^-
Spatial boundary	Γ_r
Reflective albedo	α
Angular moment	Ψ_j
Vector of angular moments	Ψ
Vector of Riemann moments	Ψ_s
Spatially discretised angular moments	$\mathbf{\Psi}$
Moment of spatially discretised angular moments	$\mathbf{\Psi}_j$
S_N quadrature set	$\{\hat{\Omega}_i, w_i\}$
Scaling functions	$\varphi_{j,k}$
Wavelet functions	$\psi_{j,m}$
Tolerance of angular adaptivity	τ

Mathematical Operators and Symbols

Gradient operator	∇
Prolongation operator	$\mathcal{Q}$
Restriction operator	$\mathcal{R}$

Acronyms

Boltzmann transport equation	BTE
Multi grid	MG
Finite element	FE
Finite element Method	FEM
Radiation non-oscillatory transport	RADIANT
Intelligent vector orthonormal replacement	IVOR
Control volume	CV
Control volume method	CVM
Discrete ordinate	DO
Discontinuous Galerkin discrete ordinate	DG-DO

The work contained chapters 3-7 of this thesis form a series of 5 papers that encompasses the work of this PhD. Their titles and the journals they have been submitted to are as follows:

1. Linear and quadrilateral octahedral wavelets on the sphere for angular discretisations of the Boltzmann transport equation [1]: Published in Annals of Nuclear Energy.

2. Linear and quadrilateral hexahedral wavelets on the sphere for angular discretisations of the Boltzmann transport equation [2]: Currently being reviewed by Nuclear Science and Engineering.

3. Chebyshev Spectral Hexahedral Wavelets on the Sphere for Angular Discretisations of the Boltzmann Transport Equation [3]: Submitted to Annals of Nuclear Energy.

4. Self adaptive spherical wavelets for angular discretisations of the Boltzmann transport equation [4]: Submitted to Nuclear Science and Engineering.

5. A hierarchical preconditioned FGMRES-based solution method for the Boltzmann transport equation [5]: Submitted to Nuclear Science and Engineering.

Chapter 1

NUMERICAL APPROXIMATION METHODS OF THE BOLTZMANN TRANSPORT EQUATION

Contents

1.1 Introduction

The ability to accurately model the transport of neutral particles has become an essential and integral part of many modern day applications. Some examples of fields which require the accurate simulation of neutral particle motion include: radiation shielding [6], radiative transfer in stellar atmospheres and clouds [7, 8], reactor simulations [9], semiconductor device design, radiation/optical oncology [10, 11, 12, 13, 14, 14] and high energy physics. The processes affecting the transport of neutral particles as they transverse through their hosting materials are described by the first order linear Boltzmann transport equation [15, 16]. The Boltzmann transport equation was derived from the study in kinetic theory of gasses where upon neutral particles were considered to behave similarly to that of a rarefied gas. The equation considers a particle's state in a 7 dimensional phase space comprising of space, angle, energy and time and conserves neutral particles by describing the individual physical processes effecting their state within this space. The processes considered are absorption, scattering and (for neutrons) production of particles by their interaction with atoms of the hosting media. Particle to particle collisions are considered inconsequential compared to particle to media interactions and thus are ignored in this transport equation.

Several analytical approaches for solving the transport equation have been developed during the first half of the 20th century. Since the 1930s and 1940s, techniques such as Weiner-Hopf [17], singular eigenfunction expansions [18] and methods using Greens functions have all been available to provide analytical solutions. However, the complexity of the 7 dimensional phase space often confines analytical methods to their use on simplified configurations, such as problem domains of infinite half spaces (for example the Milne problem). This imposes severe limitations on their use in particle transport applications and are certainly not relevant to modern day radiation engineering problems that incorporate complex multidimensional geometries. Therefore, in order to find solutions to these more general and difficult problems, numerical methods are required. In fact, the field of numerical method development now forms a substantial and important part of the research in radiation transport calculations.

It is the development of new approximation methods for the discretisation of the transport

equation's angular flux that forms the first section of research in this thesis. In particular, the research is focused on the development of discretisation methods for the directional variable of the angular flux. The essential outcome of this work is to develop robust angular discretisation methods capable of accurately representing all types of angular flux distributions. The development of new, accurate angular discretisations is particularly important to many modern day applications which often involve materials with vastly different optical properties as they induce highly complex flux distributions. One example is in radiation shielding where ducts containing voids (where no particles are absorbed) are set into highly absorbing/scattering shields in order that the system's state may be observed. The scattering material induces the flux to vary smoothly over solid angle, however, the flux within the duct can become highly directional due to particles streaming through space without colliding with the hosting material. In a similar situation, in transient reactor analysis, slug regions with optically thin cross-sections can develop within fluidised bed reactors. This is illustrated in figure 1.1 where gas bubbles combine to create large void regions within which neutrons can stream free of interaction and generate highly directional fluxes.

The implications of many problems involving material with very different optical properties are that vastly different and quickly varying angular fluxes can form over short distances. As a consequence, these distributions present numerous difficulties to the discretisation methods. Highly directional fluxes that often form in voids are notoriously difficult to approximate and require highly accurate and expensive discretisations. In fact, it is more appropriate to approximate these fluxes using adaptive technologies. In addition to this, highly directional fluxes can also cause non-physical oscillations to form over regions of angle containing high gradients (Gibbs oscillations) and these in turn may cause the angular discretisation to form spurious negative fluxes. It is important to avoid this situation, particularly for multiplying problems where particle production through fission is dependent on the current state of the flux (which therefore must be positive to be meaningful). Spurious oscillations in the scalar flux solutions (ray effects) are also a common by-product of the angular discretisation. These typically form in discrete-type discretisation solutions (such as S_N) involving problems with isolated sources and non scattering, optically thin materials.

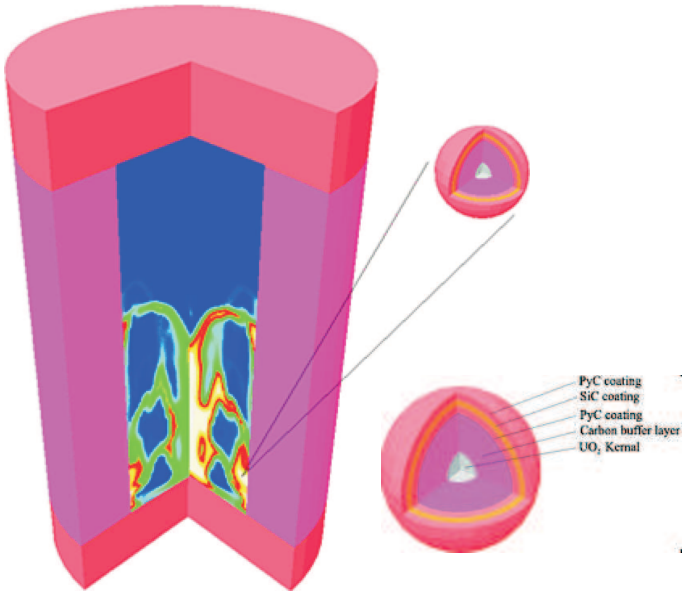

Figure 1.1: This diagram shows a conceptual fluidised bed reactor. Gas is pumped into the reactor uniformly at the bottom within which the 2mm fissile fuel particles become fluidised. As the gas (blue regions) rises through the material they coalescence and can form slugs of very low cross-sections. Neutrons then stream through these regions without collision, forming highly directional angular flux distributions.

The final part of this thesis will focus on the fast solving of the linear system of equations generated by the discretisation of the Boltzmann transport equation. This is vitally important to the progress of numerical methods as discretising all 7 dimensions of the BTE will often result in extremely large linear systems which can be immensely time consuming to solve. Specifically, a new solver based on Krylov solver technologies is developed. The solver provides fast convergence on the system's solution and reduces the computation memory costs (in comparison to standard Krylov solvers). Furthermore, in order to increase convergence rates further, a new hierarchical preconditioner is developed. The preconditioner is based on multigrid ideas utilising hierarchical angular discretisation schemes. The new solver and new preconditioner are combined to establish a fast algorithm designed specifically to solve the Boltzmann transport equation.

In order to address all issues mentioned, the new fast solver and angular discretisations are developed in order to satisfy the following:

1. To produce numerically robust angular discretisations capable of accurately representing all distributions of angular flux.
2. To mitigate or reduce ray effects and Gibbs oscillations forming in scalar flux and angular flux solutions.
3. To incorporate adaptive angular capabilities for difficult transport problems.
4. To develop a new preconditioned solver capable of increasing the efficiency of solving the discretised equations in comparison to current technologies.

The numerical methods developed in this thesis are incorporated within the radiation transport code RADIANT [19]. RADIANT employs space-time finite elements to discretise the spatial and temporal dimensions. A Streamline Upwind Petrov Galerkin method (SUPG) is incorporated within the code in order to suppress spurious oscillations forming in its numerical solutions. RADIANT has been previously compared to existing benchmarked radiation transport codes and has been shown, by numerical example, to yield a diffusion limit.

The methods of angular discretisation described here have been demonstrated for accuracy on a number of 2 dimensional, time independent, mono energetic particle transport problems. The problems have been selected so their material properties vary between being opaque and transparent in order to allow a diverse range of angular flux distributions to be tested. Benchmark solutions are provided by RADIANT using a spherical harmonic angular discretisation. This allows a consistent comparison of the angular discretisation methods to be drawn, since the remaining space and time discretisation will remain unchanged. Benchmark solutions are also provided by an even-parity [20] and discontinuous Galerkin (DG) method [21]. These are fully established radiation transport codes which therefore guarantee reliable solutions.

1.2 Numerical Methods for Solving the Boltzmann Transport Equation

This section presents a historical review of numerical methods used to model neutral particle transport for which an account of both the probabilistic and deterministic methods has been included. Since the main focus of this thesis is in the development of new angular discretisation methods, a comprehensive review of the previous angular approximation schemes is given.

1.2.1 Stochastic Methods

Monte Carlo numerical methods [22, 23] are essentially statistical modelling methods from which simulations are performed using a sequence of random numbers. For the modelling of neutral particles, hundreds of thousands of particles are traced from their birth by fission or an external source to death by capture or leakage through the problem's boundary. The path traced by a particle's movement through the problem domain is called a history. Each history tracks the absorption, scattering and fission events that a particle encounters by statistical means. Essentially, by using the probability that a particle's

interacting event will occur, random numbers between 0 and 1 can be used to determine a particle's fate at each history instance. In the event of scattering occurring, random number generators are again used to calculate the new direction of particle travel. This is based on the probability density function of its scattering angle.

A disadvantage of the Monte Carlo method is that it is difficult to predict convergence in their solutions - especially for criticality calculations. The stochastic framework also means that the method is liable to statistical errors and convergence of the solutions can not always be guaranteed. High computational cost are also a major disadvantage. This is because a large number of particle histories are required in order to gain statistically reliant results. In fact, Monte Carlo is one of the most time consuming of all methods available for particle transport modelling. The method of variance reduction has been developed to address this issue [24, 25]. This method optimises the statistics which reduces the number of particle histories required which in turn speeds up the calculation. However, for many applications the method is still considered to be insufficient.

The main advantages of using the Monte Carlo method includes the fact that it is one of the most accurate of the numerical methods available. The method avoids all discretisation errors that are associated with the decomposition of the problem domain when deterministic methods are used. In fact, the basic mechanics of Monte Carlo make it favourable for use in modelling complex multi-dimensional domains. This has led to the method being greatly applied in nuclear industry and applications of nuclear medicine. Monte Carlo modelling can also be naturally computed in parallel since particle histories act independently of each other. In summary, Monte Carlo's advantages are such that it has established itself as the main tool for particle transport modelling. It is also common to find the method being employed to provide benchmarked solutions for testing alternative numerical methods [26, 27].

1.2.2 Deterministic Spatial Discretisations

The high demand for simulating a vast range of physical phenomena, for which a solution to some fundamental partial differential equation is required, has led to an extensive range

of discretisation methods for approximating spatial variables. Of these approaches four have emerged as the most popular among researchers, namely the finite difference, finite volume, finite element and spectral methods. This section presents the main tools used to discretise the spatial domain of the Boltzmann transport equation. The methods presented are well established techniques and will have been applied to most fields of numerical modelling.

Finite Difference Method

One of the earliest discretisation schemes for approximating the BTE were designed on finite difference technologies. The finite difference method (FDM) approximates the equation over the points of a uniform grid (which is normally aligned with the Cartesian space). The approach then utilises a re-arrangement of the Taylor series and expresses the differential operators in terms of the variable's value over the grid's positions. Examples of the differencing schemes are the upwind and central differencing methods, these provide first and second order accuracy, respectively. A linear system is generated by FDM that expresses the PDE's solution over the grid positions. One of its advantages is that the treatment of the differential operators couples each grid position to their close neighbouring points only. FDM therefore produces sparse linear system which, for some equations, can be tridiagonal and symmetrical. A disadvantage of FDM is that the regular grid complicates the method's application to irregular and complex problem domains. The method's use is therefore best suited to application involving structured geometries.

Finite Element Method

The finite element method (FEM) [28, 29, 30] has emerged as one of the most powerful numerical techniques for the modelling of physical phenomena. The method first arose from the work of Rayleigh, Ritz and Galerkin who were developing new numerical solutions for problems in structural analysis. The technique proved to be successful and highly versatile, especially for modelling complicated structural domains, and it was these particularly important properties that attracted FEM to be applied in most other fields of numerical modelling.

For the modelling of neutral particle transport, finite element techniques were first de-

veloped at Los Alamos [31]. In their work they had shown how the use of higher order Lagrangians led to substantial advantages over FDM in computing the multi-group diffusion equations. From this others work soon followed. Ohnishi [32] had applied a finite element formulation to both diffusion and transport theory. Semenza too had applied the method in order to resolve the variational form of the transport equation [33]. This initial work was continued with Ackroyd [34], Kang and Hansen [35] and Pitkaranta [36], to name a few, which subsequently led to the establishment of finite element in particle transport theory. Reviews of FEM progress in this field of study, with applications to the diffusion, first order and second order equations, have been compiled by Williams and Goddard [37], Lewis [38] and Ackroyd [39].

When used to discretise the Boltzmann transport equation, the finite element method partitions the problem's spatial domain into a finite set of non overlapping elements (or sub-domains). Then, over each element, an approximation is sought using a set of interpolating polynomial functions. In most cases the approximation has linear or quadratic variation and typically uses an expansion of approximating (or trial) functions associated with collocation points positioned over the corners and mid-edges of an element. The problem is then reduced to finding the PDE's solution over the finite set of collocation points. These are computed by solving a linear system of equations which are formulated using either a weighted residual or variational method.

It is common practise to employ triangular and quadrilateral elements to discretise 2 dimensional domains. For 3 dimensional problems, tetrahedral, triangular prisms and hexahedral element are employed. In general the quadrilateral and hexahedral elements are more accurate. However, triangular and tetrahedral elements are more suitable to discretising complex geometries requiring unstructured meshes. This leads to FEM's main advantage in that complex domains are easily treated. In some instances, the method can also naturally incorporate boundary conditions.

Finite element approximations can be continuous or discontinuous over the elements' boundaries. The latter option, the Discontinuous Galerkin method (DG), is generally more accurate and stable for resolving fluxes with sharp gradients and discontinuities.

1.2.3 Deterministic Angular Discretisations

This section reviews the main numerical methods used to discretise the angular variable of the Boltzmann transport equation.

Diffusion Methods

Diffusion methods offer a simple and efficient treatment for the angular variable in particle transport modelling. Diffusion methods are low order methods that recast the transport equation into a single diffusion differential equation which is independent of angle and expressed in terms of particle density throughout space. The method began with Eddington's work in the 1940s. He had shown that if the angular flux could be represented accurately with the first two terms of a spherical harmonic expansion (that is the flux is nearly isotropic and varies smoothly over angle) the resulting equations could be reduced to a single diffusion type equation expressed only in terms of the scalar flux. Subsequent work in this area soon followed from where different assumptions were made on the angular flux distributions. Successful examples include the methods of equilibrium diffusion [40], isotropic diffusion [41] and asymptotic diffusion theory [42, 43]. In all methods, a single diffusion equation results that differed only in their diffusion coefficient.

An advantage of this approach is that the spatially discretised diffusion equations can be solved and stored in memory efficiently. The method also yields accurate solutions when the problem conforms with the assumptions of the angular flux. This usually requires the flux to have a high degree of isotropy which, in turn, makes diffusion methods suitable for many applications including reactor calculations [44] and optical imaging [11].

The methods weaknesses are that poor solutions result from difficult problems ('difficult' relating to hard angular distributions to replicate) and that breakdown occurs when resolving voids (this is due to the diffusion coefficients possessing a $\frac{1}{\sigma_t}$ term where σ_t denotes the material cross section). The methods can also violate causality. Causality states that the magnitude of the flux can not exceed the density of particles multiplied by their speed. One solution to this problem is to use flux limiting [45]. Flux limiters modify the diffusion coefficients which, in turn, manipulates the dissipation to ensure causality is maintained at all instances. Initial flux limiters were adapted in an ad-hoc manor in order to preserve

causality in the presence of large flux gradients. However, a more elegant approach was developed by Pomraning's who derived a diffusion method which had a natural built in flux limiter [46].

P_N, Simplified P_N and Solid Harmonics SHP_N

The P_N method is one of the first approximation methods used in neutral particle modelling [47, 48, 49]. The P_N method approximates the angular variable of the particle flux using expansions of spherical harmonic functions [50, 51, 52]. Spherical harmonics are a complete set of orthonormal functions that form a basis for the angular variables (μ, ω) - the variables that represent direction of particle travel on the unit sphere. The P_N approximation arises through the truncation of the spherical harmonic expansion to include all harmonics functions up to the order N. The result is a polynomial approximation of order N (over the angular variables) - where N is usually chosen to be odd. The method provides spectral accuracy and unlike the diffusion methods, it converges to the exact solution in its asymptotic limit - that is as $N \rightarrow \infty$. The spherical harmonic basis is defined as the set $\{Y^c_{l,m}, Y^s_{l,m}\}$ for $l \in \{0, 1, \ldots \infty\}$ and $m \in \{0, 1, \ldots, l\}$, where $Y^c_{l,m}$ and $Y^s_{l,m}$ denote the real valued sine and cosine spherical harmonics, respectively,

$$\begin{aligned} Y^c_{l,m} &= P^m_l(\mu)\ \cos(m\omega), \\ Y^s_{l,m} &= P^m_l(\mu)\ \sin(m\omega). \end{aligned} \tag{1.1}$$

The functions P^m_l define the associated Legendre polynomials [15].

The advantages of P_N is that it is generally accurate for most transport problems. For problem domains that involve significant scattering, and where the angular flux possesses a fair degree of isotropy, low order expansions are sufficient for accurate approximations. The method is also rotationally invariant. This is particularly important as ray effects are guaranteed not to form in the scalar flux solution. P_N also has the ability to treat anisotropic scattering kernels exactly. In addition to this, when P_N is applied to the second order transport equations the odd moments can be eliminated to form the reduced set of even parity equations.

The disadvantages of P_N include the problem of vacuum boundaries being difficult to maintain due to the continuous variation of the approximated flux. There is also the

problem of sharp gradients in the scalar flux solution being smoothed out by low order expansions. However, the more significant disadvantages are that the multidimensional P_N equation sets are complicated, expensive to solve and grow rapidly in the order of $(N+1)^2$. This has been a huge draw back for the method's progress in the transport approximations. Even despite being one of the earliest methods developed, the high computational requirements has pushed P_N beyond use in most practical applications.

In contrast to this, the one dimensional P_N equations, which are derived by an expansion of Legendre polynomials [15], are simple and grow at the much slower rate of just $(N+1)$. It was this reduction in complexity that attracted Gelbard to the idea of expanding the planar equation sets into a multi-dimensional framework as a method of avoiding the complications of P_N. He first achieved his goal in the 1960s [53] when he proposed the simplified spherical harmonic method (SP_N). His method was to generalise the set of $(N+1)$ planar equations by the replacement of the spatial derivatives with multidimensional Laplacian and divergence operators. The initial test calculations on one dimensional cylindrical [54] and spherical [55] geometries yielded impressive results.

By its construction, the SP_N advantage is that the equations are simple and small in number. In addition to this, the method retains its rotational invariance, therefore ensuring the solutions do not contain ray effects. Unfortunately the method's mathematical foundation was weak for general domains and Gelbards ad-hoc derivation was never fully accepted by the research community. In addition to this, the asymptotic limit of SP_N did not converge to the equation's exact solution (except on a few specific cases). For these reasons the research in this field stagnated over the next two decades. It was not until the early 1980's that interest in SP_N began to re-emerge. In 1981, Lemanska [56] reported huge improvements in accuracy on diffusion theory using the low order SP_3 and SP_5 equations. Smith [57, 58] then independently derived the SP_N equations and again reported huge increases in accuracy over the diffusion theory applied to two dimensional lattice and core applications. As interest grew, SP_N was applied to multidimensional applications of shielding, core and lattice problems [59] and again similar improvements on the diffusion equations were reported. The general pattern showed SP_N was capable of providing significant improvements on the diffusion theory (or P_1), especially when the diffusion (P_1) was the leading term. This view was strengthened when Gamino [60, 61]

reported that over 80% of the error of the P_1 approximation could be corrected with a low order SP_N solution.

The main breakthrough came when papers by Larsen, McGhee and Morel [62, 63, 64] gave a mathematical basis for the asymptotic limit for the SP_N equations. Using both a variation analysis and asymptotic limit argument, they were able to show that SP_N was a high order approximation of the transport equation for problems where the P_1 expansion was the leading term. Geometries with this attribute have the tendency of being optically thick, having low absorption cross-sections and relatively high scattering cross-sections which are not highly forward peaked. Essentially, SP_N is a higher order approximation to the diffusion theory. Pomraming was also able to show that for domains with an infinite homogeneous medium, SP_N solutions were guaranteed to converge on the exact solution.

Another simplification of the P_N approximation is the solid harmonic method (SPH_N) proposed by Ackroyd in the 1960s [65, 66, 67]. Ackroyd's aim was to relax the constraints of the unit vector denoting direction and the inconveniences caused by its directional components. The transport equation was generalised to incorporate a directional vector of arbitrary length and expanded in angle using solid harmonics. The result was a set of multi-dimensional first order differential equations that had strong similarities to the planar P_N equations. Subsequent work followed which recast the equations into second order differential form. Elimination of the odd moments then provided a system of diffusion-like equations that showed a strong resemblance to those of SP_N.

Discrete Ordinates

The discrete ordinate method (S_N) is one of the most popular and simple techniques used in radiation transport approximations. The discrete ordinate method discretises the angular variable into a finite set of rays or ordinates, $\{\hat{\Omega}_i\}$ for $i \in \{1, 2, \ldots, \mathcal{M}\}$, and particles are permitted to travel only in these directions. The directions are then coupled with a set of weights to form a quadrature set on the sphere, $\{\hat{\Omega}_i, w_i\}$ for $i \in \{1, 2, \ldots, \mathcal{M}\}$, and the transport equation's integrals are numerically approximated using these S_N quadrature rules. The set of $\mathcal{M}$ equations follows from expressing the transport equation at each of the directions $\hat{\Omega}_i$, for $i \in \{1, 2, \ldots, \mathcal{M}\}$.

The S_N method was proposed by Carlson [68] in the 1950s and quickly became popular for transport calculations. A reason for its popularity was that its simple treatment of the angular variable meant that the equations sets were only coupled through scattering. Therefore, pure absorbing problems would generate a set of $\mathcal{M}$ independent hyperbolic equations that are solved for each direction. In addition to this, bare surface boundary conditions are easily resolved by simply determining the ordinates directing particles into the system and setting the associated fluxes to zero.

The S_N method does have many disadvantages, the most notable of which are the presence of ray effects in their solutions. Ray effects [69, 70] arise in multidimensional solutions as a direct consequence of the discretisation limiting the angular flux to travel in a finite set of directions. Ray effects appear as unphysical oscillations in scalar flux solutions and are at their most damaging when a problem contains isolates sources within a rarefied, non scattering media. Various methods and techniques have been proposed in order to suppress the oscillations caused by ray effects from destroying S_N solutions. One technique is to ensure enough directions are used so that source regions are directly connected to all elements of the spatial discretisation. Another technique is to select a quadrature set that is more rotationally invariant, thus less susceptible to oscillating solutions. A third, more elegant approach is to recast the S_N equations into a set with close similarities to the P_N equations (that are free from ray effects) [71, 72]. This is achieved by adding a fictitious source term to the transport equation to induce direction coupling through the streaming operator. In addition to ray effects, another S_N disadvantage is that it generates a fully coupled system of equations for scattering problems. This has obvious disadvantages for the memory and CPU resources.

Various S_N direction and weighting quadrature sets have been developed. For 1-D calculations the direction and weights are normally chosen to correspond to the roots of the $N-1$ Legendre polynomial [15]. In this instance it can be shown that the S_N and P_{N-1} approximations are equivalent [15]. For multi-dimensional discretisations, common quadrature sets are usually equal weight numerical integration schemes. Two examples are the Kershaw $\&$ Harte and rectangular partitioning methods [73], an example of these is presented in figure 1.2. The diagrams illustrate the method's discretisation by a partitioning on the sphere. The centre of each patch denotes an S_N direction while its area

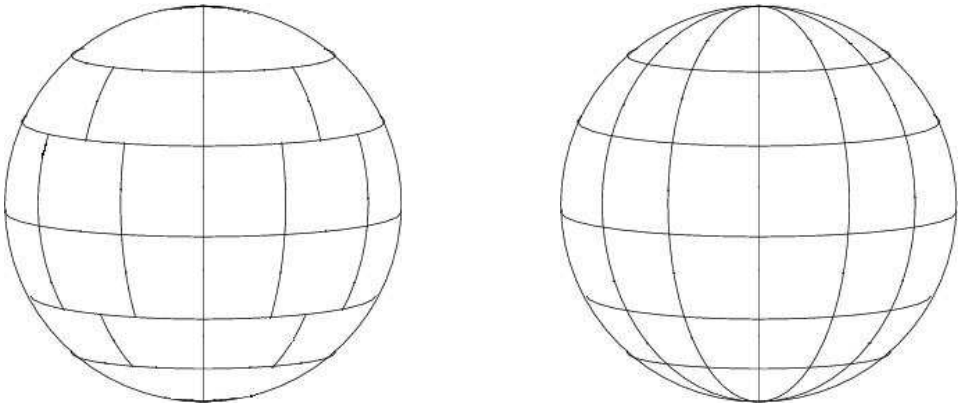

(a) (b)

Figure 1.2: The equal weight discrete ordinate discretisation of the sphere. (a) The Kershaw & Harte method (S_6). (b) The rectangular partitioning method. The centre of the patches represent a direction of S_N. The patch's area represent the direction's corresponding weight in the S_N quadrature rule.

corresponds to the weight. Other multi-dimensional quadrature schemes include uniform weight [74], uniform Gauss-weight (using the tensor product of two 1-D Gauss quadrature rules) [75] and Chebyshev-Legendre [76] schemes.

Walsh Functions

Walsh functions [77, 72, 78] are well established in numerical analysis for expanding and approximating functions on the real line. Walsh functions form an orthonormal and complete set of hierarchical basis functions over the intervals $x \in [-1, 1)$. The hierarchical levels, $j \in \{0, 1, \ldots, \infty\}$, form piece-wise constant approximation over the intervals $(2^j + k, 2^{j+1} + k]$ for $k \in \{0, 1, \ldots 2^j\}$. The functions are defined as:

$$Wal(2i+p, x) = (-1)^{[i/2]+p} \left\{ Wal(i, 2(x+0.5)) + (-1)^{i+p}\, Wal(i, 2(x-0.5)) \right\} \quad (1.2)$$

with an initial Walsh function defined as,

$$Wal(0, x) = \begin{cases} 1 \text{ for } -1 \leq x < 1 \\ 0 \text{ otherwise,} \end{cases} \quad (1.3)$$

and the first four Walsh functions are presented in figure 1.3.

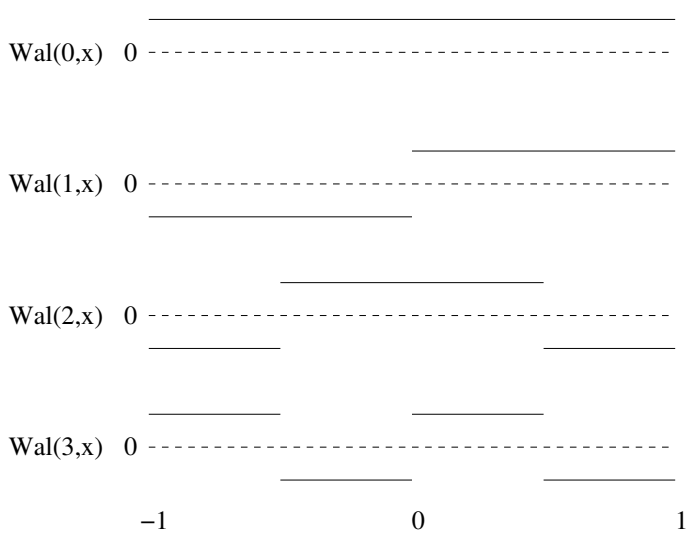

Figure 1.3: This diagram presents the first 4 Walsh functions defined in equation 1.2.

A tensor product of two Walsh expansions are used to discretise the angular variable of the Boltzmann transport equation. The two expansions are used to discretise angular variable over the azimuth and polar angles separately. This results in an approximation that is piece wise constant over a rectangular partitioning of the sphere which is similar to the second equal weight S_N partitioning shown in figure 1.2. The Walsh approximation has an advantage over S_N in that the light coupling of the angular variables through the streaming term serves to reduce oscillating solutions caused by ray-effects [78]. In addition to this, the discretisation of scattering terms does not result in fully coupled systems (as what happens with S_N). Furthermore, the two separate expansions allow the angular discretisation to vary its resolution across the polar and azimuth directions. This allows high resolution to be applied to the directional variables that are important - for instance the azimuth angle is more important in 2-D calculations [79]. The disadvantage of the method is that ray effects do occur in the solutions. The method is also of low order accuracy.

1.2.4 Solutions of the Boltzmann Transport Equation

This section describes the two main methods for solving the discretised Boltzmann transport equation. A popular technique for reducing spurious oscillation that result from the discretisation are also discussed.

Variational Techniques

Variational techniques are employed when the solution to the differential equation under consideration is equivalent to the optimisation of a different differential equation [80]. The variational method functions by substituting an approximation into this alternative equation (called the functional). The functional is then optimised by differentiating it with respect to each expansion coefficient and solving the resulting linear system with all terms equated to zero. One advantage of using variational techniques is that the functional is often easier to handle as it is usually of lower differential order than the original differential equation. Another advantage is that the method can be used to prove the existence of a solution.

The variational method has been extensively applied to obtain solutions for the second order even parity equations. Examples of these include Lillie and Robinson work on discrete ordinate, finite element solutions of the even and odd parity equation [81] and the spherical harmonic, finite element solutions of multigroup problems developed by Oliveria [82]. A full demonstration of the variational approach's use in transport calculations is also given in [49].

Weighted Residual Methods

Residual methods are techniques that find solutions to the discretised equations by the minimisation of its residual term. The residual is defined as the error that results from the substitution of the approximation into the transport equation. Obviously, the expansion of trial functions will not normally be able to provide a solution that results in a zero residual over whole phase space. Therefore, as an alternative, the residual is required to vanish in some average sense. One residual approach is the collocation method. This technique selects a finite number of collocation points at which the residual is required

to vanish. Other popular techniques include the sub-domain collocation and least squares methods. However, the most popular of the residual methods is the weighted residual Galerkin technique. In the Galerkin method the approximated equation is multiplied by a set of weighting functions and integrated over its phase space. This results in a set of weak form equations that form to build a linear system. This is then solved to find the expansion coefficients that satisfies the Galerkin weighted conditions. Numerous weighting techniques exist. One approach is to perform the weighting using the same trial functions as used in the approximation - the Bubnov Galerkin method. Alternatively one could use a separate set of trial functions to carry out the weighting - the Petrov Galerkin method.

Stabilization Methods

For many numerical techniques, solving PDEs in which convection dominates can often result in spurious Gibbs oscillations forming in their approximations. Typically, under high convection conditions, sub optimal numerical schemes produce solutions that are under diffusive and contain spurious oscillations over regions of high gradients. For many applications in transport modelling, the oscillations can be highly damaging and may also cause negative angular and scalar fluxes to form. In these situations it becomes necessary to mitigate the oscillations and negative solutions using a suitable stabilization method.

Numerous techniques have been developed for numerical stability in finite element and finite difference solutions. The Petrov Galerkin methods are one such successful example. In Petrov-Galerkin methods, additional diffusion is introduced into the solution in order to compensate the under dissipative approximation of the original discretisation. This is achieved by pre-multiplying the transport equation by the Petrov Galerkin term, this is essentially a contribution of the equations convection operator. Then, when combined with the equation's own advection, the SUPG term yields an additional diffusion component. Selecting the optimal dissipative contribution of the stabilizing term remains an active area research. Various techniques have been derived that perform well under specific material regimes. For example, the self adjoint angular flux method (SAAF) provides stability for optically thick materials. The Streamline Upwind Petrov Galerkin method [83] (SUPG), in which diffusion is added in the solution's stream line direction only, is more applicable for optically thin regimes. Methods combining both approaches in order

to gain optimal stability over all regimes have also been developed [84, 19]. In addition to this, non-linear methods can also be applied to SUPG in order to mitigate oscillations that do not propagate in the stream line direction - such examples include discontinuity or shock capturing methods [84, 19].

1.3 Scope of Thesis

This thesis is concerned with approximating the solution of the time independent, mono energetic, first order linear Boltzmann transport equation. The first aim of this thesis is to produce new robust numerical methods for discretising the transport equation's angular variable. This work is based on the application of spherical wavelets for which three schemes have been developed. Two are based on low order wavelet discretisations and the other is an approximation method using spectrally accurate wavelets. The reasons for investigating spherical wavelet discretisations are due to their particular properties. Apart from having already shown their powerful ability in function approximation, their properties of compact support and hierarchical expansions makes the application of anisotropic adaptive resolution in the angular phase-space possible. This leads into the next topic of the thesis - to develop and apply a method for self adaptive angular discretisations of the Boltzmann transport equation. The final section of this thesis concentrates on the efficient solving of the space/angle discretised equations. In this work a new variation of FGMRES is developed. The solver is also combined with a new preconditioner based on hierarchical angular multigrid ideas in order to increase efficiency further.

In chapter 2 the transport of neutral particles is described and the first order linear Boltzmann transport equation is presented. The angular discretised transport equations are then derived for an arbitrary angular discretisation. The full treatment of the scattering term has been included where the scattering cross-sections are assumed anisotropic and expanded using Legendre polynomials. A Riemann method for treating bare surface boundary conditions is also presented. Chapters 3 and 4 develop the two low order spherical wavelet discretisations based on second generation wavelets. The full development is presented and tests are performed on a number of 2-D problems domains. Chapter 5 presents the

spectral wavelet discretisation based on Chebyshev collocation point polynomials. Again tests are conducted on a number of 2-D problems domains. Chapter 6 develops a new method for self adaptive angular discretisations using the spherical wavelet bases developed in the previous chapters. Numerical examples are presented showing the successful functioning of the adaptive angular resolution using the two low order wavelet bases. In chapter 7 the new solver and preconditioner are developed. The solver is based on Arnoldi Krylov search space ideas and is an extension of the FGMRES algorithm. The preconditioner uses the hierarchy of the angular discretisations which naturally lends itself to a hierarchical or multi-grid solution method. This allows the preconditioner complete freedom to select and concentrate on resolving the angular moments that require the highest effort. Chapter 8 presents the conclusions of this thesis and a section detailing possible future work.

Chapter 2

THE BOLTZMANN TRANSPORT EQUATION

Contents

2.1 Introduction

The Boltzmann Transport Equation (BTE) [49, 15, 16] mathematically describes the transport of neutral particles, such as neutrons and photons, and details the interaction

processes between the particles and their hosting materials. The transport equation was derived by postulating some basic concepts, or assumptions, concerning the directional flow of radiation and how its orientation and energies are influenced when contact is made with the atoms of the surrounding matter. The assumptions conform exceptionally well to the natural motion of neutral particles and naturally lead to the notation of material cross-sections and secondary particles emitted through the processes of scattering and fission reactions.

This chapter presents the foundations of particle transport simulation that are relevant for the work of this thesis to proceed. Section 2.1.1 presents the assumptions of neutral particle travel and describes the terminology of material cross sections that describe the events of absorption, scattering and fission. Section 2.1.2 then presents the linear Boltzmann transport equation (as well as its alternative forms) and describes boundary conditions common to practical applications. Section 2.2 derives the angular discretised Boltzmann transport equations. The equations are based on using arbitrary expansions of angular basis functions to approximate the transport equation's angular dependence. The treatment of the scattering term using an expansions of Legendre polynomials are also fully described. Section 2.3 presents an SUPG method and a finite element discretisation of the spatial variables. A treatment of bare surface boundary conditions using a Riemann method is presented. This method is illustrated on 1-D equations where the Riemann method is used to approximate fluxes at the surfaces of a control volume discretisation.

2.1.1 Neutral Particle Motion

For the derivation of the Boltzmann Transport Equation the following presumptions on neutral particle propagation are:

1. Particles are considered as points.
2. Particles travel in straight lines between collisions with the surrounding matter.
3. Interaction between particles may be ignored.

4. Particles emitted from collisions with the surrounding matter are instantaneous, except in the case of delayed neutrons in fission reactions.

5. The material-particle interaction properties are assumed to be known.

A consequence of assumptions 2 and 3 is that a particle's energy and direction can only change through collisions with the atoms of the hosting materials. The types of interactions that occur on collision are absorption, scattering and, for neutrons, fission. In the absorption process a particle is captured by an atom of the hosting media and, in a mathematical point of view, cease to exist thereafter. In a scattering interaction, a particle is deflected by contact with a hosting material atom which results in a change to its direction and energy state. The two types of scattering processes are elastic and inelastic. Elastic scattering is governed by the conservation of momentum and kinetic energy where as inelastic scattering results from the loss of a particle's kinetic energy that occurs due to the increase in the atoms energy state. Finally, in a fission reaction, a neutron is captured (absorbed) by the nucleus of a hosting fissile atom. An instantaneous or delayed emission of one of more neutrons then follows this event.

The rates of which all three processes occur are governed by the material cross-sections which are known from the outset by assumption 5. The cross-sections, denoted by σ, are macroscopic quantities. They have the dimensions cm^{-1} and represent the probability that a particle will undergo a collisional process having travelled a distance of dr. For example, the probability of a particle being captured when travelling a distance of dr through a material with an absorbing cross-section σ_a, is given by,

$$\mathrm{absorption\ \ probability} = \sigma_a \, dr.$$

In a similar manor the probabilities of scattering and fission processes occurring are described by the cross-sections σ_s and σ_f respectively. The likelihood of any collision occurring is denoted by the material's total cross-section σ_t. This is defined as the sum of all interaction cross-sections,

$$\sigma_t \quad = \quad \sigma_a \quad + \quad \sigma_s,$$

where the absorption cross-section σ_a is the summation of the capture σ_c and fission σ_f cross-sections,

$$\sigma_a \quad = \quad \sigma_c \quad + \quad \sigma_s.$$

The total cross-section's reciprocal also defines the materials mean free path. This is the average distance a particle can expect to travel before undergoing a collision event.

In this thesis only materials with isotropic cross-sections are considered. This requires the probability of a particle's interaction occurrence to be independent of its initial direction. For the vast majority of applications this property conforms well with the behaviour of materials. However, in some instances the application of anisotropic materials are important [85]. Some examples include applications of modelling radionuclides through fractured rock [86], calculations of low energy neutron propagation through crystalline media and near-inferred propagation through living tissue.

2.1.2 The Boltzmann Transport Equation

This section presents the first order, time dependent linear Boltzmann transport equation which describes the transport of neutral particles based on the assumptions stated in section 2.1.1,

$$\begin{aligned} \frac{1}{v}\frac{\partial}{\partial t}\psi(r,\hat{\Omega},E,t) \;&+\; \hat{\Omega}.\nabla\psi(r,\hat{\Omega},E,t) \;+\; \sigma_t(r,E)\psi(r,\hat{\Omega},E,t) \;= \\ q_{ex}(r,\hat{\Omega},E,t) \;&+\; q_s(r,\hat{\Omega},E,t) \;+\; q_f(r,\hat{\Omega},E,t). \end{aligned} \tag{2.1}$$

The field variable of the Boltzmann transport equation is the angular flux $\psi(r,\hat{\Omega},E,t)$. This variable exists in a 7 dimensional phase space that is made up of three spatial, two angle, an energy and a time dimension. The angular flux represents the number of neutral particles that occupy the spatial region r, that are travelling in the direction $\hat{\Omega}$, that have the energy E at the time instance t, multiplied by their velocity v.

The Boltzmann transport is essentially a balance equation that conserves neutral particles during their propagation through a material. Each component of equation (2.1) mathematically represents the individual mechanisms influencing particles over the phase space

instance $(r, \hat{\Omega}, E, t)$. The first term of equation (2.1) represents the rate of change of angular flux with respect to time t. In order to ensure the number of particles balance, this term equates to the number of particles lost through streaming, the number of particles lost through absorption and the number of particles gained through emission during the short time period $t + dt$. The second and third terms of equation (2.1) describe the net loss of particles due to streaming and capture (for all interactions), respectively. The three terms on the right side of equation (2.1), q_{ex}, q_s and q_f, denote the emission of angular flux coming from external sources, scattering and fission reactions, respectively. The external source $q_{ex}(r, \hat{\Omega}, E, t)$ is an explicitly known function of the 7 dimensional phase space. The emission of particles through scattering and fission are implicit functions which are dependent on the current state of angular flux.

The discharge of scattered particles in direction Ω with energy E is calculated in a statistical average framework. Using the cross-section $\sigma_s(r, \hat{\Omega}' \rightarrow \hat{\Omega}, E' \rightarrow E, t)$ to denote the probably density function of a particle, initially travelling in direction $\hat{\Omega}'$ with energy state E', to be scattered into the direction $\hat{\Omega}$ with energy E, the scattering emission can be calculated by integrating the cross-section and angular flux $\psi(r, \hat{\Omega}', E', t)$ over the angle and energy phase space (Ω', E'),

$$q_s(r, \hat{\Omega}, E, t) = \int_{E'} \int_{4\pi} \sigma_s(r, \hat{\Omega}' \rightarrow \hat{\Omega}, E' \rightarrow E, t)\psi(r, \hat{\Omega}', E', t)d\hat{\Omega}'dE'. \tag{2.2}$$

The emission of particles through fission (where delayed neutrons neglected) is also calculated in a statistical framework. For this a term $\chi(E)$ is used to denote the probability density function of a fission neutron having an energy E. A second term $\nu(E')$ denotes the mean number of neutrons produced in a fission reaction initiated by the absorption of a neutron with energy E'. The fission cross section $\sigma_f(r, E')$ is used to represent the probability of a neutron with the energy E' being captured by fission. The assumption made on these terms is that the discharge of fission is assumed to have isotropic distribution - that is, the neutrons released have no bias to any direction. The output of fission particles can then be calculated by multiplying the probability and cross section terms with the scalar flux and integrating over the energy variable E',

$$q_f(r, \hat{\Omega}, E, t) = \chi(E) \int_0^{\infty} \nu(E')\sigma_f(r, E')\phi(r, E', t)dE'. \tag{2.3}$$

The scalar flux is defined as the angular flux integrated over angle $\hat{\Omega}$.

$$\phi(r, E, t) = \int d\Omega \psi(r, \hat{\Omega}, E, t) \tag{2.4}$$

This has an important physical meaning as it defines the mean density of particles at position r with energy E at time t. Two other important quantities are the radiative flux and pressure due to the radiation flow. The radiative flux is a three element vector measuring the flow of particles across a unit area of surface at position r. This is defined by,

$$\underline{\phi}(r, E, t) = \int d\Omega \, \underline{\hat{\Omega}} \, \psi(r, \hat{\Omega}, E, t), \tag{2.5}$$

where $\underline{\hat{\Omega}}$ is the unit vector defining the direction Ω in terms of the spatial coordinates. The pressure due to the radiation flow is represented by the 9 element matrix, again defined in terms of the unit vector $\underline{\hat{\Omega}}$ and the angular flux,

$$\hat{\phi}(r, E, t) = \int d\Omega \, \underline{\hat{\Omega}} \, \underline{\hat{\Omega}} \, \psi(r, \hat{\Omega}, E, t). \tag{2.6}$$

Alternative forms of the BTE

It is important to mention alternative versions of the Boltzmann transport equation as established codes that solve these variations are frequently employed to provide benchmark solutions. One variation is the even-parity equation which can be constructed directly from the first order equation (2.1) or by means of a variational approach [34]. The method defines the even and odd parity variables as,

$$\psi^{+}(\hat{\Omega}) = \frac{1}{2}(\psi(\hat{\Omega}) + \psi(-\hat{\Omega})),$$
$$\psi^{-}(\hat{\Omega}) = \frac{1}{2}(\psi(\hat{\Omega}) - \psi(-\hat{\Omega})),$$

which replace the angular flux within the transport equation. Cancellation of the odd-parity variables, the second order equation involving the even-parity moments results,

$$-\hat{\Omega}.\nabla \frac{1}{\sigma_t} \hat{\Omega}.\nabla \psi^{+}(r, \hat{\Omega}, E, t) + \sigma_t \psi^{+}(r, \hat{\Omega}, E, t)$$
$$= q_{ex}(r, \hat{\Omega}, E, t) + q_s(r, \hat{\Omega}, E, t) \; + \; q_f(r, \hat{\Omega}, E, t). \tag{2.7}$$

An advantage of this equation is that it is self-adjoint, which enables fast and efficient algorithms to be applied for solving the discretised equations. Furthermore, a spherical

harmonic approximation of the angular flux results in the elimination of the odd harmonic moments, which reduces the number of angular unknowns which in turn increases the efficiency of solving the discretised equations. The main disadvantage of the even-parity method is the presence of the $\frac{1}{\sigma_t}$ term inside the second order differential operator. This prevents the equation's use in solving particle transport through voids. In this situation it becomes necessary to add small cross sections to void regions in order to prevent the equations breaking down. However, one must be sure that these small variations do not have large influences on the solution.

Another alternative version of the transport equation is in its integral form. This version is again constructed from the first order equation, the details for this can be found in [15] and [16]. The time independent integral form transport equation is given by,

$$\psi(r,\Omega) = \Gamma(r_s,\Omega)\ \exp(-\int_0^{|r-r_s|} ds^{''}\sigma_t(r-s^{''}\hat{\Omega})) + \int_0^{\infty} ds^{'} q(r-s^{'}\hat{\Omega},\hat{\Omega})\exp(-\int_0^{s'} ds^{''}\sigma_t(r-s^{''}\hat{\Omega})). \tag{2.8}$$

The point denoted by r_s is the position on the boundary which is in line with the opposite direction of Ω from position r. The term $\Gamma(r_s,\Omega)$ denotes the boundary condition at position r_s, boundary conditions will be discussed in the next section. The term s denotes a spatial distance and q is the sum of all flux emitted through external sources, scattering and fission. For pure absorbing problems, $\sigma_s = \sigma_f = 0$, equation (2.8) provides an analytical solution of the Boltzmann transport equation.

2.1.3 Boundary Conditions

In order to solve the transport equation, certain conditions at the system's boundary must be specified. The initial state of the angular flux at the first time instance (t=0) must be known explicitly for all space, angle and energy dimensions,

$$\psi(r,\hat{\Omega},E,0) = \Psi(r,\hat{\Omega},E). \tag{2.9}$$

Furthermore, conditions at the system's spatial boundary Γ_r are required for all time instances. More precisely, it is the incoming angular flux at the spatial boundary that is

required to satisfy certain boundary conditions. That is, the angular flux is required for,

$$\psi(r, \hat{\Omega}, E, t), \quad \forall \quad E,\ t,\ r \in \Gamma_r,\ \hat{\Omega}.n < 0, \tag{2.10}$$

where n denotes the outward normal vector perpendicular to the surface boundary at the point $r \in \Gamma_r$.

In practical applications of particle transport modelling a number of common boundary conditions are applied over a domain's spatial edge. Explicit conditions include surface sources at the spatial boundaries,

$$\psi(r, \hat{\Omega}, E, t) = q_{ex}(r, \hat{\Omega}, E, t) \tag{2.11}$$

for $r \in \Gamma_r$ and $\hat{\Omega}.n < 0$. For the case where the source is zero, i.e. there is no incoming flux at the boundary, the boundary is said to be a bare surface or vacuum boundary. Common implicit boundary conditions include the albedo reflecting condition. This defines the incoming angular flux to be the reflected image of the outgoing flux multiplied by an albedo constant $\alpha(E)$,

$$\psi(r, \hat{\Omega}, E, t) = \alpha(E)\ \psi(r, \hat{\Omega}^*, E, t) \quad r \in \Gamma_r,\ \hat{\Omega}.n < 0. \tag{2.12}$$

The angle $\hat{\Omega}^*$ represents the outgoing reflecting angle to $\hat{\Omega}$ with respect to the boundary at the point $r \in \Gamma_r$. This can be calculated by satisfying the conditions,

$$\begin{aligned} \hat{\Omega}.n &= -\hat{\Omega}^*.n, \\ (\hat{\Omega} \times \hat{\Omega}^*).n &= 0. \end{aligned} \tag{2.13}$$

In the case of $\alpha(E) = 1$, the boundary condition (2.12) gives a perfect reflecting surface.

A second implicit surface condition is the White boundary condition. The White condition requires all particles passing through the boundary at point $r \in \Gamma_r$ to be distributed back into the system isotropically. The incoming distribution of flux for this surface condition is represented by,

$$\psi(r, \hat{\Omega}, E, t) = \frac{1}{4} \int_{n.\hat{\Omega}'>0} d\hat{\Omega}'\ n.\hat{\Omega}'\ \psi(r, \hat{\Omega}', E, t). \tag{2.14}$$

2.1.4 The BTE Considered in the Thesis

This thesis considers the mono energetic, non multiplying linear Boltzmann transport equation. In order to reduce the mathematical notation, the obsolete energy term E and fission source q_f are not considered any further from here. The thesis assumes the scattering cross-sections are anisotropic. That is, the probability of a scattered particle's direction is only dependent on the scattered angle $\mu_0 = \hat{\Omega}'.\hat{\Omega}$. This in turn leads to the following simplification of the scattering emission of equation (2.2),

$$q_s(r, \hat{\Omega}, t) = \int_{4\pi} \sigma_s(r, \mu_0, t)\ \psi(r, \hat{\Omega}', t)\ d\hat{\Omega}'. \tag{2.15}$$

The final form of the Boltzmann transport equation that is considered in this thesis is given as:

$$\begin{aligned} \frac{1}{v}\frac{\partial}{\partial t}\psi(r, \hat{\Omega}, t) \ &+ \ \hat{\Omega}.\nabla\psi(r, \hat{\Omega}, t) \ + \ \sigma_t(r)\psi(r, \hat{\Omega}, t) \ = \\ q_{ex}(r, \hat{\Omega}, t) \ &+ \ \int_{4\pi} \sigma_s(r, \mu_0, t)\ \psi(r, \hat{\Omega}', t)\ d\hat{\Omega}'. \end{aligned} \tag{2.16}$$

2.2 The Angularly Discretised Transport Equation

This section derives the angular discretised Boltzmann transport equations. The work in the following chapter of this thesis will begin from this point. The following sections describe the full procedure for discretising the angular dimension using an arbitrary angular approximation scheme. This includes a complete description of treating the scattering cross-sections using expansions of Legendre polynomials. This section begins from the Boltzmann transport equation (2.16)

2.2.1 The Non-Scattering System of Equations

This section considers the transport equation (2.16) without the scattering term q_s. In order to approximate the angular variable and derive the angular discretised transport equations, the angular dimensions of (2.16) are expanded using some arbitrary set of

angular basis functions. That is, the angular flux $\psi(r, \hat{\Omega}, t)$ is approximated using an expansion of the form,

$$\psi(r, \hat{\Omega}, t) \approx \sum_{j=1}^{\mathcal{M}} \Psi_j(r, t)\mathcal{G}_j(\hat{\Omega}), \tag{2.17}$$

where $\mathcal{G}_j(\hat{\Omega})$ represent the expansion's angular basis functions and $\Psi_j(r, t)$ are the corresponding expansion coefficients (in this case they are referred to as the angular moments). The approximation (2.17) is inserted into the transport equation (2.16) and a system of $\mathcal{M}$ equations are generated using a weighted residual method. This requires equation (2.16) to be multiplied by a weighting function $\mathcal{G}'_i(\hat{\Omega})$, for $i \in \{1, \ldots, \mathcal{M}\}$, and integrated over the angular variable Ω. For the work in this thesis the Bobnov Galerkin method is applied. This method uses the angular basis functions in the expansion (2.17) as the weighting functions, i.e. $\mathcal{G}' = \mathcal{G}$. This results in,

$$\int d\Omega\, \mathcal{G}_i(\hat{\Omega}) \Big\{ (\frac{1}{v}\frac{\partial}{\partial t} + \hat{\Omega}.\nabla + \sigma_T(r, E)) \sum_{j=1}^{\mathcal{M}} \Psi_j(r, t)\mathcal{G}_j(\hat{\Omega}) - q_{ex}(r, \hat{\Omega}, E, t)\Big\} = 0 \qquad \text{for } j = 1, \ldots, \mathcal{M}. \tag{2.18}$$

From this point, the vector r will be described by the Cartesian coordinates $\{x, y, z\}$ and the angular vector Ω is represented by the variables μ , the cosine of the polar angle θ measured with respect to the spatially fixed z axis, and the azimuthal angle ω, the angle describing the direction of Ω in the xy plane - see figure 2.1. This set of spatial and angle coordinates allows the streaming operator of (2.18) to be defined specifically as,

$$\hat{\Omega}.\nabla = (\hat{\Omega}_x, \hat{\Omega}_y, \hat{\Omega}_z).(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}) = \hat{\Omega}_x\frac{\partial}{\partial x} + \hat{\Omega}_y\frac{\partial}{\partial y} + \hat{\Omega}_z\frac{\partial}{\partial z}, \tag{2.19}$$

where $\hat{\Omega}_x, \hat{\Omega}_y$ and $\hat{\Omega}_z$ are components of the unit vector describing the direction of Ω in Cartesian space. These components are expressed in terms of μ and ω by,

$$\begin{aligned} \hat{\Omega}_x &= (1-\mu^2)^{\frac{1}{2}}\cos\omega, \\ \hat{\Omega}_y &= (1-\mu^2)^{\frac{1}{2}}\sin\omega, \\ \hat{\Omega}_z &= \mu. \end{aligned} \tag{2.20}$$

The angularly discretised transport equations are derived by inserting (2.19) and (2.20) into (2.18), and grouping the removal and scattering operators (which is discussed in the

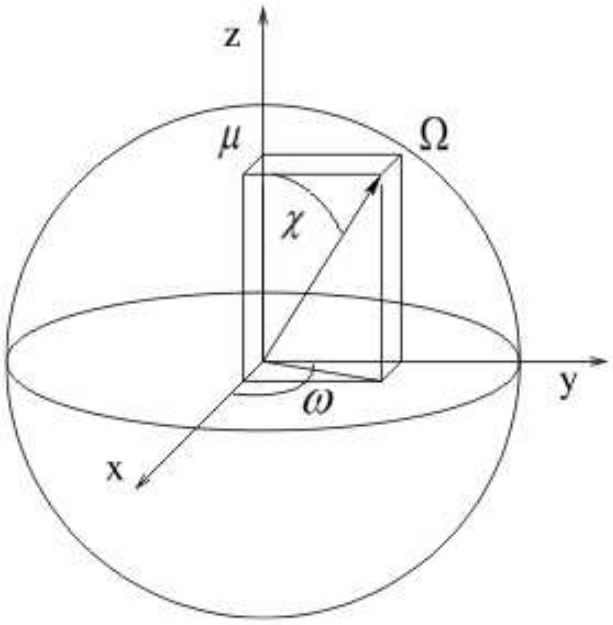

Figure 2.1: This diagram illustrates the representation of particle travel on the unit sphere. The direction of particle travel is represented by Ω. This variable is constructed from the two variables ω and χ. These represent the azimuth and co-latitude angles, respectively. In many cases it is convenient to use the cosine of χ, this is represented by μ.

next section) together.

$$\frac{1}{v}A\frac{\partial\Psi(r,t)}{\partial t}+A_x\frac{\partial\Psi(r,t)}{\partial x}+A_y\frac{\partial\Psi(r,t)}{\partial y}+A_z\frac{\partial\Psi(r,t)}{\partial z}+\mathcal{H}\Psi(r,t)=\mathcal{S}(r,t) \quad (2.21)$$

The variable $\Psi(r,t)$ is a vector containing the $\mathcal{M}$ angular moments which form the approximation of the angular flux at position r and time t. The variable $\mathcal{S}(r,t)$ is a vector of size $\mathcal{M}$ which forms the angularly discretised external source over position r and at time t. A is the $\mathcal{M}\times\mathcal{M}$ angular mass matrix, A_x, A_y and A_z are the $\mathcal{M}\times\mathcal{M}$ angular Jacobin matrices and $\mathcal{H}$ is the $\mathcal{M}\times\mathcal{M}$ scattering removal operator matrix. The elements of the above matrices and source vectors are specified as,

$$\begin{aligned}
A_{i,j} &= \int d\Omega\, G_i(\hat{\Omega})\, G_j(\hat{\Omega}) &&= \int d\mu \int d\omega\, G_i(\mu,\omega)\, G_j(\mu,\omega),\\
A_{xi,j} &= \int d\Omega\, \hat{\Omega}_x\, G_i(\hat{\Omega})\, G_j(\hat{\Omega}) &&= \int d\mu \int d\omega\, (1-\mu^2)^{\frac{1}{2}}\cos(\omega)\, G_i(\mu,\omega)\, G_j(\mu,\omega),\\
A_{yi,j} &= \int d\Omega\, \hat{\Omega}_y\, G_i(\hat{\Omega})\, G_j(\hat{\Omega}) &&= \int d\mu \int d\omega\, (1-\mu^2)^{\frac{1}{2}}\sin(\omega)\, G_i(\mu,\omega) G_j(\mu,\omega),\\
A_{zi,j} &= \int d\Omega\, \hat{\Omega}_z\, G_i(\hat{\Omega})\, G_j(\hat{\Omega}) &&= \int d\mu \int d\omega\, \mu\, G_i(\mu,\omega) G_j(\mu,\omega),\\
\mathcal{S}_j(r,t) &= \int d\Omega\, q_{ex}(r,\hat{\Omega},t)\, G_j(\hat{\Omega}) &&= \int d\mu \int d\omega\, q_{ex}(\mu,\omega)\, G_j(\mu,\omega). \qquad (2.22)
\end{aligned}$$

The elements of the scattering/removal matrix $\mathcal{H}$ are defined in the next section.

2.2.2 Discretisation of the Scattering Term

In order to discretise anisotropic scattering cross sections, it is customary to expand the probability density function of the scattered angle μ_0, used in equation (2.15), in terms of the orthogonal Legendre polynomials [15]. That is, the scattering cross section $\sigma_s(\mu_0)$ is expanded by,

$$\sigma_s(r,\hat{\Omega}.\hat{\Omega}') = \sigma_s(r,\mu_0) = \sum_{l=0}^{\infty}(2l+1)\,\sigma_{sl}(r)\,P_l(\mu_0), \quad (2.23)$$

where P_l is the l^{th} order Legendre polynomial and $\sigma_{sl}(r)$ are their corresponding expansion coefficients. Due to the orthogonality of the Legendre polynomial basis, the expansion coefficients σ_{sl} can be computed by multiplying the cross section with the l^{th} order Legendre polynomial and integrating over the scattering angle μ_0.

$$\sigma_{sl} = \int_{-1}^{1} \frac{d\mu_0}{2} \, \sigma_s \, (\mu_0) P_l(\mu_0) \tag{2.24}$$

Inserting (2.23) into (2.15) gives the following expression for the emission density of scattered particles,

$$q_s(r, \hat{\Omega}, t) = \sum_{l=0}^{\infty} (2l+1)\sigma_{sl} \int d\Omega' \, P_l(\hat{\Omega}.\hat{\Omega}')\psi(r, \hat{\Omega}', t). \tag{2.25}$$

This is then simplified using the addition theorem [15] which expands each Legendre polynomial in terms of products of spherical harmonic functions,

$$P_l(\hat{\Omega}.\hat{\Omega}') = \sum_{m=-l}^{l} Y^*_{l,m}(\hat{\Omega}) Y_{l,m}(\hat{\Omega}'). \tag{2.26}$$

Complex spherical harmonics [15] are defined as,

$$Y_{l,m}(\hat{\Omega}) = P_l^m(\mu) e^{i\omega m}, \tag{2.27}$$

with $Y^*_{l,m}(\hat{\Omega})$ denoting their complex conjugates and P_l^m represent the associated Legendre polynomials [15].

The purpose of this is that the directions before and after scattering, $\hat{\Omega}$ and $\hat{\Omega}'$, can be separated and treated individually. Combining (2.25) and (2.26) yields the following expression for particles emitted through scattering,

$$q_s(r, \hat{\Omega}, t) = \sum_{l=0}^{M} \sum_{m=-l}^{l} Y^*_{l,m}(\hat{\Omega}) \sigma_{sl}(r) \int d\Omega' Y_{l,m}(\hat{\Omega}') \psi(r, \hat{\Omega}', t), \tag{2.28}$$

which is now in a suitable form for the angular dimension to be discretised.

Conforming with the previous section, the Bobnov Galerkin weighted residual method is employed to discretise angular variable. Equation (2.28) is pre-multiplied by the angular basis function $\mathcal{G}_i$, for $i \in \{1, 2, \ldots, \mathcal{M}\}$ and integrated over angle Ω,

$$\int d\hat{\Omega} \, G_j(\hat{\Omega}) \sum_{l=0}^{M} \sum_{m=-l}^{l} \, Y^*_{l,m}(\hat{\Omega}) \, \sigma_{sl}(r) \int d\Omega' \, Y_{l,m}(\hat{\Omega}') \, \psi(r, \hat{\Omega}', t),$$
$$\forall i \in \{1, 2, \ldots, \mathcal{M}\}, \tag{2.29}$$

which is then arranged into the form,

$$\int d\hat{\Omega}\, G_j(\hat{\Omega}) \sum_{l=0}^{M} Y^*_{l,0}(\hat{\Omega})\, \sigma_{sl}(r) \int d\Omega'\, Y_{l,0}(\hat{\Omega}')\, \psi(r, \hat{\Omega}', t) \quad +$$
$$\int d\hat{\Omega} G_j(\hat{\Omega}) \sum_{l=0}^{M} \sum_{m=-l_{m\neq 0}}^{l} Y^*_{l,m}(\hat{\Omega}) \sigma_{sl}(r) \int d\Omega' Y_{l,m}(\hat{\Omega}') \psi(r, \hat{\Omega}', t),$$
$$\forall i \in \{1, 2, \ldots, \mathcal{M}\}. \tag{2.30}$$

This expression can be simplified using the following property of the complex spherical harmonics,

$$Y^*_{l,-m}(\hat{\Omega}) = (-1)^m Y_{l,m}(\hat{\Omega}). \tag{2.31}$$

Allowing $\Phi^m_l(r,t)$ to represent the integrals of equation (2.30),

$$\Phi^m_l(r,t) = \int d\Omega'\, Y_{l,m}(\hat{\Omega}')\, \psi(r, \hat{\Omega}', t), \tag{2.32}$$

the sum $Y^*_{l,m}(\hat{\Omega})\Phi^m_l + Y_{l,-m}(\hat{\Omega})\Phi^{-m}_l$ can be expressed as,

$$\begin{aligned}
&Y^*_{l,m}(\hat{\Omega})\, \Phi^m_l \;+\; Y_{l,-m}(\hat{\Omega})\, \Phi^{-m}_l \\
&= Y^*_{l,m}(\hat{\Omega}) \int d\hat{\Omega}' Y_{l,m}(\hat{\Omega}')\, \Psi(\hat{\Omega}') + (-1)^m\, Y_{l,m}(\hat{\Omega}) \int d\hat{\Omega}' (-1)^m\, Y^*_{l,m}(\hat{\Omega})\, \Psi(\hat{\Omega}') \\
&= \; [Y^e_{l,m}(\hat{\Omega}) - i\, Y^o_{l,m}(\hat{\Omega})] \int [Y^e_{l,m}(\hat{\Omega}') + i\, Y^o_{l,m}(\hat{\Omega}')\, \Psi(\hat{\Omega}')]\, d\hat{\Omega}' \\
&+ \; [Y^e_{l,m}(\hat{\Omega}) + i\, Y^o_{l,m}(\hat{\Omega})] \int [Y^e_{l,m}(\hat{\Omega}') + i\, Y^o_{l,m}(\hat{\Omega}')\, \Psi(\hat{\Omega}')]\, d\hat{\Omega}' \\
&= \int Y^e_{l,m}(\hat{\Omega}')\, \Psi(\hat{\Omega}')\, d\hat{\Omega}'\, [Y^e_{l,m}(\hat{\Omega}) - iY^o_{l,m}(\hat{\Omega}) + Y^e_{l,m}(\hat{\Omega}) + i\, Y^o_{l,m}(\hat{\Omega})] \\
&+ i \int Y^o_{l,m}(\hat{\Omega}')\Psi(\hat{\Omega}')\, d\hat{\Omega}'\, [Y^e_{l,m}(\hat{\Omega}) - i\, Y^o_{l,m}(\hat{\Omega}) - Y^e_{l,m}(\hat{\Omega}) - i\, Y^o_{l,m}(\hat{\Omega})] \\
&= 2\, Y^e_{l,m}(\hat{\Omega}) \int Y^e_{l,m}(\hat{\Omega}')\, \Psi(\hat{\Omega}')\, d\hat{\Omega}' + 2\, Y^o_{l,m}(\hat{\Omega}) \int Y^o_{l,m}(\hat{\Omega}')\, \Psi(\hat{\Omega}')\, d\hat{\Omega}'
\end{aligned} \tag{2.33}$$

where the functions $Y^e_{l,m}(\hat{\Omega})$ and $Y^o_{l,m}(\hat{\Omega})$ represent the real and complex parts of the spherical harmonic function $Y_{l,m}(\hat{\Omega})$,

$$\begin{aligned}
Y^e_{l,m}(\hat{\Omega}) &= \sqrt{\frac{(2l+1)(l-m)!}{(l+m)!}} P^m_l(\mu)\cos(m\omega), \\
Y^o_{l,m}(\hat{\Omega}) &= \sqrt{\frac{(2l+1)(l-m)!}{(l+m)!}} P^m_l(\mu)\sin(m\omega),
\end{aligned} \tag{2.34}$$

where

$$Y_{l,m}(\hat{\Omega}) = Y^e_{l,m}(\hat{\Omega}) + i\, Y^o_{l,m}(\hat{\Omega}).$$

Using the result of (2.33) the scattering contribution of equation (2.30) is expressed in terms of the spherical harmonics' real and complex parts, Y^e and Y^o.

$$\int d\hat{\Omega}\, G_j(\hat{\Omega}) \sum_{l=0}^{M} \sigma_{sl}(r)\, Y^e_{l,0}(\hat{\Omega}) \int d\Omega'\, Y^e_{l,0}(\hat{\Omega}')\, \psi(r,\hat{\Omega}',t) + \sum_{l=1}^{\infty} \sigma_{sl}(r) \sum_{m=1}^{l} [2\, Y^e_{l,m}(\hat{\Omega}) \int d\Omega' Y^e_{l,m}(\hat{\Omega}')\, \psi(r,\hat{\Omega}',t) + 2\, Y^o_{l,m}(\hat{\Omega}) \int d\Omega' Y^o_{l,m}(\hat{\Omega}')\, \psi(r,\hat{\Omega}',t)] \tag{2.35}$$

Then, by inserting the angular approximation (2.17) into (2.35) and defining the variables,

$$\alpha^{e,l,m}_j = \int d\Omega\, G_j(\hat{\Omega})\, Y^e_{l,m}(\hat{\Omega}),$$
$$\alpha^{o,l,m}_j = \int d\Omega\, G_j(\hat{\Omega})\, Y^o_{l,m}(\hat{\Omega}), \tag{2.36}$$

the equation set (2.35) can be expressed as,

$$q_s(r,\hat{\Omega},t) = \sum_{l=0}^{\infty} \sigma_{sl}(r)\, \alpha^{e,l,0}_j \sum_{j=1}^{\mathcal{M}} \alpha^{e,l,0}_i\, \Psi_i(r,t) + 2\sum_{l=1}^{\infty} \sigma_{sl}(r) \sum_{m=1}^{l} [\alpha^{e,l,m}_j \sum_{i=1}^{\mathcal{M}} \alpha^{e,l,m}_i \Psi_i(r,t) + \alpha^{o,l,m}_j \sum_{j=1}^{\mathcal{M}} \alpha^{o,l,m}_i \Psi_i(r,t)], \tag{2.37}$$

for $i \in \{1, 2 \ldots, \mathcal{M}\}$. Upon the re-arrangement of the above equation, that is, by pulling angular coefficients outside the brackets, the weighted scattering term is given by,

$$\sum_{j=1}^{\mathcal{M}} \Psi_i(r,t)[\sum_{l=0}^{\infty} \sigma_{sl}(r)\, \alpha^{e,l,0}_j \alpha^{e,l,0}_i + 2\sum_{l=1}^{\infty} \sigma_{sl}(r) \sum_{m=1}^{l} [\alpha^{e,l,m}_j\, \alpha^{e,l,m}_i + \alpha^{o,l,m}_j\, \alpha^{o,l,m}_i]], \quad \forall i \in \{1,2,\ldots,\mathcal{M}\}. \tag{2.38}$$

The equation (2.38) is arranged into a matrix vector form, the vector being the set of angular moments. Therefore, the scattering contribution to the (i,j) element of $\mathcal{H}$ in (2.21) is calculated to be the expression within the brackets. The elements of $\mathcal{H}$ in (2.21)

are now defined to be,

$$\mathcal{H}_{i,j} = \int d\mu \int d\omega \; G_i(\mu,\omega)\, G_j(\mu,\omega) \; - \sum_{l=0}^{\infty} \sigma_{sl}(r)\, \alpha_j^{e,l,0} \alpha_i^{e,l,0} + 2 \sum_{l=1}^{\infty} \sigma_{sl}(r) \sum_{m=1}^{l} (\alpha_j^{e,l,m}\, \alpha_i^{e,l,m} + \alpha_j^{o,l,m}\, \alpha_i^{o,l,m}), \quad (2.39)$$

where the integrals on the right side account for the total cross-section removal term.

A Summary of the Angular Discretised Equations

The full angular discretisation of the non-multiplying linear Boltzmann transport equation using an arbitrary angular approximation (2.17) has been presented. For all the angular discretisation methods discussed in the following chapters, the objective is to derive the system of equations (2.21) for which the matrices and vectors are defined by (2.22) and (2.39). At this point the angular flux is fully described by the coefficients $\Psi_j(r,t)$ for $j \in \{1, 2, \ldots, \mathcal{M}\}$, and forms the starting point for developing the time and spatial discretisation methods.

2.3 Spatial Discretisation Methods

This section briefly overviews the spatial treatments of the Boltzmann transport equation that are used in the numerical examples of this thesis. The aim here is not to provide a complete description of the spatial treatment, as this will be covered in the following chapters, but rather to present the spatial discretised transport equations at which the point of resolving boundary conditions becomes necessary. This will then be followed by a comprehensive description for treating bare surface boundary conditions using a Riemann method. In order to simplify the equations the time dimension has been neglected as it has no implication on the treatment of the boundary condition.

2.3.1 SUPG Stabilization

Discretisation of the spatial variables for the multi dimensional numerical examples in this thesis use both the SUPG finite element and control volume finite element (CVFEM) methods. In both approaches the SUPG stabilization method is employed in order to reduce un-physical spatial oscillations forming over regions with large gradients [84, 19]. The SUPG method is described in more detail in the later chapters of this thesis. Here it is just briefly introduced in order that the spatially discretised equations can be derived. The method of SUPG pre-multiplies the angularly discretised transport equation (2.21) by the SUPG term $(I - \boldsymbol{A} \cdot \nabla P)$, where $\boldsymbol{A}$ denotes the vector of matrices containing the angular Jacobians, $\boldsymbol{A} = (A_x, A_y, A_z)$, and P is the stabilizing matrix. The SUPG term is designed to add diffusion to the transport approximations in their streamline direction only . The elements of the matrix P are selected in order to add the correct diffusion for a stable, non-oscillatory solution (which will be addressed in a later chapter).

2.3.2 CVFEM & FEM Discretisations

In this section a brief outline of the control volume finite element method (CVFEM) and finite element method (FEM) is presented. In these two approaches the resulting discretised equations follow similar paths in their development, and vary only when a different set of weighting functions are used. It is assumed that the angular flux spatial domain has been discretised using some arbitrary finite element mesh.

From the SUPG angular discretised equations, the equation sets are once again pre-multiplied by a set of weighting functions and then integrated over the spatial domain. For the FEM, the weighting functions are the same functions used to discretise and approximate the spatial variable of the angular flux. For the CVFEM, the weighting functions are defined using a control volume discretisation. The control volumes are centred over the nodes used in the finite element discretisation and a formed by a Voronoi decomposition [87] of the spatial domain. The weighting functions M_i, for $i \in \{1, 2, \ldots, \mathcal{N}\}$ have the value 1 over their associated control volume CV_i and are 0 else where.

For both methods, the resulting set of spatially discretised equations are,

$$\int_V \boldsymbol{M}_i(r)(\mathcal{H}\Psi(r,t) - \mathcal{S}(r,t))dV + \int_{\Gamma_r} \boldsymbol{M}_i(r)(\boldsymbol{A}\cdot n)\Psi(r,t)) = 0, \quad \forall\, i \ \in \{1,\ldots\mathcal{N}\},$$

for CVFEM (see [88]) and,

$$\begin{aligned} -\int_V (\boldsymbol{A}\cdot\nabla\boldsymbol{N}_i(r))\Psi(r)dV \quad &+ \quad \int_V \boldsymbol{N}_i(r)(H(r)\Psi(r) - S(r))dV \\ +\int_V \boldsymbol{A}\cdot\nabla\boldsymbol{N}_i(r)P\mathcal{R}(\Psi(r))dV \quad &+ \quad \int_{\Gamma_r} \boldsymbol{N}_i(r)(\boldsymbol{A}\cdot\boldsymbol{n})(\Psi(r))d\Gamma_r, \\ &\forall\, i \ \in \{1,\ldots\mathcal{N}\}, \end{aligned}$$

for FEM (see [83]). The term $\boldsymbol{N}_i$ denotes the $\mathcal{M}\times\mathcal{M}$ diagonal matrix containing the finite element function associated with node i of the finite element mesh. Similarly, the term $\boldsymbol{M}_i$ denotes the $\mathcal{M}\times\mathcal{M}$ diagonal matrix containing the control volume function associated with node i of the finite element mesh. The term n denotes the outward pointing normal vector on the boundary Γ_r.

2.3.3 Riemann Method for Bare Surfaces

In both the CVFEM and FEM methods, similar surface integrals appear in the discretised transport equations as a result of applying Greens theorem to the volume integrals,

$$\int_{\Gamma_r} \boldsymbol{M}_i(r)(\boldsymbol{A}\cdot n)\Psi(r,t)), \qquad \forall\, i \ \in \{1,\ldots\mathcal{N}\}, \tag{2.40}$$

and

$$\int_{\Gamma_r} \boldsymbol{N}_i(r)(\boldsymbol{A}\cdot\boldsymbol{n})(\Psi(r))d\Gamma_r, \qquad \forall\, i \ \in \{1,\ldots\mathcal{N}\}. \tag{2.41}$$

In order to resolve these integrals correctly, the angular flux must also satisfy the system's boundary conditions. For the numerical examples in this thesis, vacuum boundary conditions are resolved using a Riemann approach [89, 19]. The basic mechanics of the Riemann method is to exploit the eigenstructure of the angular Jacobian matrices in order to decouple the angular moments and calculate the incoming and outgoing information passing through the system's boundary. In order to illustrate Riemann method a 1-D control volume discretisation method is presented (with no SUPG) in which the Riemann

method is employed to approximate the incoming and outgoing information through each of the control volume's surface.

1-D CV Riemann Decomposition

For this demonstration the steady state, 1-D angular discretised Boltzmann transport equation is considered without a scattering/removal operator. The spatial domain is discretised using a set of $\mathcal{N}$ control volumes (for simplicity a regular spaced grid is considered) and the angular flux is spatially expanded using,

$$\Psi(r) \approx \boldsymbol{\Psi}(r) = \sum_{j=1}^{\mathcal{N}} \boldsymbol{\Psi}_j \boldsymbol{M}_j(r), \tag{2.42}$$

where $\boldsymbol{M}_j(r)$ denote the control volume functions (which have the value 1 over their respective control volume CV_i and 0 else where) and $\boldsymbol{\Psi}_j$ are the expansion's coefficients. Applying the standard Bubnov Galerkin weighted residual method, the angular discretised equations are pre-multiplied by each control volume and integrated over space (which is just the x variable in this case),

$$\int_{CV_i} M_i A_x \frac{\partial \Psi(x)}{\partial x} dx = \int_{CV_i} M_i S(x), \qquad i \in \{1, 2, \ldots, \mathcal{N}\}. \tag{2.43}$$

The angular Jacobian A_x is then expressed in terms of its eigenstructure by,

$$A_x = R_x \Lambda_x R_x^{-1}, \tag{2.44}$$

where R_x is an $\mathcal{M} \times \mathcal{M}$ matrix formed from the augmentation of the $\mathcal{M}$ columns of right eigen vectors of A_x. The matrix Λ_x is an $\mathcal{M} \times \mathcal{M}$ diagonal matrix containing the eigen values of A_x along its diagonal elements. Due to the symmetry of A_x the matrix of left eigenvectors can be represented by the inverse of R_x in the decomposition (2.44).

The method then maps the angular moments $\Psi(x)$ onto a Riemann space using a mapping matrix defined by R_x,

$$\Psi(x) = R_x \Psi_x(x). \tag{2.45}$$

Inserting the mapping (2.45) into (2.43) and pre-multiplying by R_x^{-1} gives,

$$\int_{CV_i} M_i \Lambda_x \frac{\partial \Psi_x(x)}{\partial x}\, dx = R_x^{-1} \int_{CV_i} M_i S(x), \quad i \in \{1, 2, \ldots, \mathcal{N}\}. \tag{2.46}$$

Applying Green's theorem to the streaming term reduces the order of the differential operator and simplifies the volume integral to one over the surface of the control volume,

$$\int_{\Gamma_{CV_i}} M_i n_x \Lambda_x \Psi_x(x) d\Gamma = R_x^{-1} \int_V M_i S(x,t), \quad i \in \{1,2,\ldots,\mathcal{N}\}, \tag{2.47}$$

where the term n_x denotes the outward direction to the control volume's surface. The implication of this approach is that in equation (2.47), the Riemann variables are decoupled through the streaming term. This is crucial as each variable may now be treated separately, effectively reducing the problem to a set of independent advection equations. This can be used to greatly reduce the complexity of approximating the moment variables at the boundary, which is required by the integral in equation (2.47). As the direction of advection for each Riemann variable Ψ_{xi} is defined solely by the sign of its corresponding eigenvalue $\Lambda_{xi,i}$, it is easy to determine whether a Riemann variable is streaming information into or out of of a control volume at its surface. Therefore, an estimation of the flux at the control volume's surface can be calculated using upwind approximations in Riemann space. This approach has essentially divided the surface integral of equation (2.47) into two integrals representing information entering and leaving the control volume at its surface.

$$\int_{\Gamma_{CV_i}} M_i n_x \Lambda_x \Psi_x(x) d\Gamma = \int_{\Gamma_{CV_i}} M_i n_x \Lambda_x \Psi_x^{in}(x) d\Gamma + \int_{\Gamma_{CV_i}} M_i n_x \Lambda_x \Psi_x^{out}(x) d\Gamma \tag{2.48}$$

The Riemann variables that stream out of a control volume are approximated at a boundary by taking the values of the control volume's Riemann variables. The Riemann variables that stream into the control volume are estimated at the boundary by taking the value of the Riemann variables associated with the neighbouring control volume. The integrals of equation (2.47) are then calculated in Riemann space and mapped back to the angular moment space using the mapping (2.45). The diagram presented in figure 2.2 gives an illustration of the Riemann method.

Riemann Solution of the Multi-Dimensional Surface Integrals

In order to resolve bare surfaces and satisfy the surface integrals (2.40) and (2.41), the Riemann method applies the ideas presented in the 1-D control volume example. The method once again uses the eigenstructure of the matrix $A_s = (\boldsymbol{A} \cdot n)$,

$$A_s = R_s \Lambda_s R_s^{-1}, \tag{2.49}$$

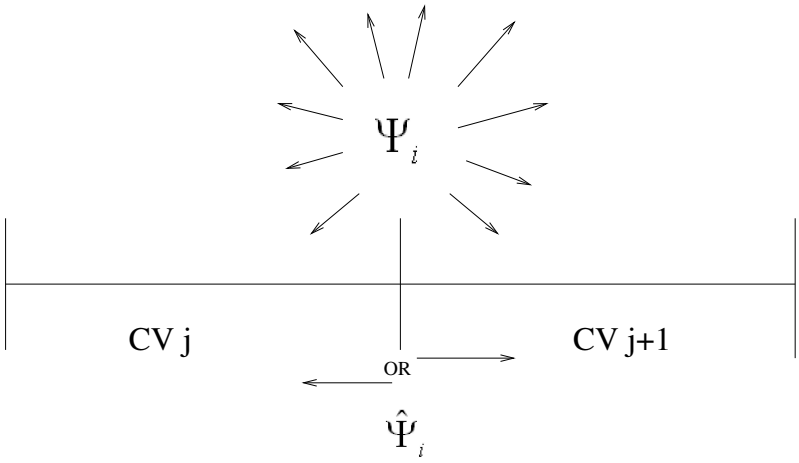

Figure 2.2: This diagram illustrates how the Riemann method works when estimating the flux at the edge of a control volume for a 1-D problem. The approximation of each angular moment is normally quite difficult due to the fact that its associated angular function can often direct flux over large areas of angle - as illustrated by the arrows of the variable Ψ_i. Since these directions of flow will often transport particles into and out of a control volume an up-wind estimation of the variable is difficult. However, by mapping the angular moments to a Riemann space, each Riemann variable, denoted by $\hat{\Psi}_i$, directs information in one specific direction. A consequence of this is that the Riemann information is being directed either in or out of the control volume and so an upwind estimation of this variable is straight forward. Estimations of the variables at the surface of a control volume are therefore made in Riemann space and the resulting approximations then mapped back to the proper angular moment variables.

where, as before, the matrices R_s and Λ_s denote the right eigenvectors and eigenvalues of A_s, respectively. Again, a mapping between the angular moments and Riemann variables is defined in terms of the right eigenvectors,

$$\Psi(x) = R_s \Psi_s(x). \tag{2.50}$$

On substitution of expressions (2.49) and (2.50) into the surface integrals (2.40) and (2.41), the following results,

$$\int_\Gamma \boldsymbol{M}_i A_s \Psi(r,t) d\Gamma = \int_\Gamma \boldsymbol{M}_i R_s \Lambda_s R_s^{-1} R_s \Psi_s(r,t) d\Gamma = R_s \int_\Gamma \boldsymbol{M}_i \Lambda_s \Psi_s(r,t) d\Gamma, \tag{2.51}$$

and,

$$\int_\Gamma \boldsymbol{N}_i A_s \Psi(r,t) d\Gamma = \int_\Gamma \boldsymbol{N}_i R_s \Lambda_s R_s^{-1} R_s \Psi_s(r,t) d\Gamma = R_s \int_\Gamma \boldsymbol{N}_i \Lambda_s \Psi_s(r,t) d\Gamma. \tag{2.52}$$

The surface integration can now be resolved in Riemann space by determining the Riemann variables that direct information in and out of the system through the boundary. In order to satisfy bare surfaces the incoming Riemann variables are set to zero. Out going variables are approximated at the boundary using either upwind approximations or finite element interpolation. It should also be mentioned that when integrating over the surface, Gauss points located on the internal of the surface are used to evaluate the integral. Therefore, as elements with flat edges are used, the direction n remains in a constant direction over all quadrature points on a surface - also note that the directions at the nodes will not need to be considered. This therefore keeps the angular Riemann decomposition independent of space, thus the mapping between the angular and Riemann space remains constant on an element's surface.

The importance of this Riemann method for resolving boundary conditions is that it is defined using an arbitrary angular discretisation. This therefore allows the same code to be used with any approximation of the angular flux. All that is required is a re-definition of the angular matrices in equation (2.21).

Chapter 3

LINEAR AND QUADRATIC OCTAHEDRAL WAVELETS ON THE SPHERE FOR ANGULAR DISCRETISATIONS OF THE BOLTZMANN TRANSPORT EQUATION

Contents

This chapter develops two new wavelet bases for discretising the angular term of the first order Boltzmann transport equation. The wavelets functions are based on Sweldens second generation wavelets [90] that are generated using the lifting procedure [91]. In this chapter the wavelets are built on an octahedral representation of the sphere and the angular flux approximation takes the form of finite element linear and quadratic representations. The spherical wavelets are similar to the functions developed in [92] and [93]. However, in this derivation the bases apply a new amendment for mitigating the ray effect type inaccuracies that can dominate 1 and 2 dimensional solutions. A demonstration of the new angular discretisation techniques are performed on a number of 1 and 2 dimensional problem domains. These numerical examples aim to demonstrate these new methods capabilities to accurately represent the angular flux and to determine their susceptibility to ray effects forming in their solutions. Benchmark solutions are provided by the conventional S_N and P_N discretisations that are used to solve both the first order and even parity equations.

3.1 Introduction

In recent years wavelets [94, 95, 96] have proved themselves to be a versatile tool in many fields of mathematics, physics, statistics, engineering and computational sciences. The key to their success lies with their power to efficiently represent general functions and data sets. In effect, wavelets have the ability to efficiently encompass and evaluate data utilising only a small number of coefficients.

Wavelet theory first attracted the attention of several researchers such as Morlet, Grossmann, Mayer and Daubechies [95] in the early 1980s. Their work had a tremendous impact on techniques in signal processing and image and data compression for which wavelet analysis has since become both an established and popular method. This popularity was based on the method's functions used in the representation and analysis of signals, since, unlike the established method of Fourier analysis which uses bases of global sines and cosines, the wavelet functions, or wavelets, were localised in both frequency and time. The implications of this property not only lead to a naturally preserved space-time localised wavelet transform (analogous to the Fourier transform) but the compact wavelets were easier to implement than the global functions of their Fourier counterparts. Furthermore, the invention of the Multi-resolution analysis [94] enabled a wavelet reconstructing of a signal to be performed in layers of increasing resolution. Each incrementation of the reconstruction involves supplementing the existing expansion with a set of wavelets giving the required increase in accuracy. In effect, this resulted in a hierarchical wavelet basis constructing the local profile of a signal in terms of its information at various resolutions. This immediately lead to the prospects of an adaptive reconstruction of the signal in which high resolution analysis would be confined to the problem's regions in which it was required. Together, these factors were to become the greatest strength of wavelet analysis in numerical applications.

For the applications of numerical analysis classical wavelets constitute a basis for $L^2(R)$. They take the form of translations and dilation of a single 'mother wavelet' $\varphi(x)$ and posses the desirable properties of orthogonality, compact support and zero integrals. They are formally defined in terms of their mother wavelet,

$$\varphi_{j,k}(x) = 2^{j/2}\varphi(2^{j/2}x - k), \qquad \forall j, k \in \mathbf{Z}, \tag{3.1}$$

for which the set $\{\varphi_{j,k}, j, k \in Z\}$ forms the dense basis of $L^2(R)$. The number of wavelets has continued to grow since their formalism, each developed to possess the properties suitable for the specific problem at hand, and thus have created a large family of wavelet methods. Common examples include piece-wise constant Haar functions, spline polynomials, Legendre, Daubechies and Gaussian wavelets [94].

As previously mentioned, wavelets play a pivotal role in image analysis. Wavelets have

been applied to applications including edge detection [97, 98], de-noising or filtering [99, 100] and, most significantly, image compression [101, 102, 103]. Wavelet based 2-D image compression techniques commonly employ a tensor product extension of (3.1) to two variables, referred to as a separable transform, although non-separable transforms have been developed [99]. Examples presented in [104] and [105] show wavelet compression technologies to be comparable or even to exceed the performance of other leading compression techniques. These papers include the well established wavelet compression algorithms EZW [104] and SPIHT [105].

Following their early success in signal analysis and data compression, wavelets attracted the attention of researches from various other fields. One such field was in numerical techniques for solving partial differential equations (PDE) where the prospect of an efficient and adaptive algorithm was foreseen as an attractive rival to the traditional techniques of finite element, finite difference and spectral methods. Wavelets have since been applied across the whole spectrum of PDEs. An example of these include their application in solving Hyperbolic [106], Elliptic, Stokes/Helmoltz [107] and Burgers equations [108, 109]. An extensive reference list of their application in PDEs up to 2001 can be found in [110].

In almost all classical wavelet based approximations of multidimensional PDEs, the separate spatial directions are individually expanded, from which the full domain representation follows from their tensor product. In most of these applications wavelets expand all spatial dimensions, however, they have also been mixed with other approximation techniques with considerable success. A mixed Spectral/Wavelet method was proposed in [107] for the solution of the 2-D stokes equation for which the flow had periodicity in one direction. A Daubechies wavelet expansion was used for representing the periodic flow, while a spectral Chebyshev expansion was employed for the representation of the non periodic direction. Numerical examples demonstrated the method to be both stable and spectrally accurate.

Wavelet based discretisation techniques fall into either of two categories: Galerkin [106, 107] and collocation methods [111, 108, 109, 107]. In the Galerkin method the PDE is projected onto the wavelet space through its multiplication with the wavelet basis and integration over the problem domain. This scheme in general leads to better pre-conditioning

and conforms well with the wavelets theoretical analysis. In the collocation scheme the wavelets are associated with unique collocation points. The PDE is then expressed at these collocation points in order to obtain the fully discretised system of equations. The method has the advantage that it can easily treat any non linear terms, there also exists a number of techniques to compute derivatives. The method also offers a potential route for resolving boundary conditions. However boundary conditions are in general a cause for some concern. This is due to wavelets being constructed over the infinite real line which therefore possess no natural method of satisfying boundaries. One successful solution was to develop wavelets and wavelet transforms on an interval [112, 113, 114]. Another technique was to completely reconstruct the classical wavelet theory, stripping away the translations and dilations of a mother wavelet and replacing it with a new method of wavelet construction. So drastic were these changes from the classical analysis that this new breed of wavelets became known as Second Generation Wavelets [115].

Second Generation Wavelets were founded by Swelden in the mid 1990s through his work in restructuring wavelet analysis into a more general framework [90, 116]. The approach superseded Fourier analysis as the tool for wavelet construction, which had been used for all wavelet development with the exception of interpolating functions developed by Donoho [117] and Harten [118]. Instead, wavelets were developed through a simple but effective method called the lifting scheme [90, 119]. The basic idea was to start with very simple wavelet functions that held the essential properties to co-exist with a given multiresolution analysis. The lifting scheme then operated on the wavelets working them up to functions with particular useful properties. This alternative approach did come at a cost as some of the important properties of the classical wavelets were inevitably lost. However, despite this disadvantage, the method did provide the necessary freedom for satisfying boundary conditions. Furthermore the functions could also retain their properties such as compact support and zero integrals through the weight settings of the lifting coefficients. If these positive benefits were not enough, the huge benefits of the second generation wavelets was that they were constructed on the spatial domain. This enabled an easy adaption over complex and irregular geometries and even their extension to other more exotic topologies from R and R^2.

Research in computer graphics has received a considerable amount of activity through

the use of both classical and second generation wavelets [120, 96]. Examples of classical wavelet applications include B-spline and Daubechies wavelets in surface and volume illumination computations [121, 122], surface modelling [123] and animation [124]. The second generation wavelets made their largest contributions in the field of three dimensional surface imaging [96]. This realisation was first set out in the work of Lounsbery [125] who developed a wavelet basis that could construct complex, three dimensional topologies comprising of triangular elements or faces. These arbitrary structures would be generated from recursive subdivisions of a much simpler triangular structure - for example a tetrahedron. Each subdivision involved splitting each face into 4 new triangular faces which were then manipulated by the wavelets to pose a closer resemblance with their intended shape - for example a human head. The approach gave the user a control over the resolution of the whole image by increasing or decreasing the number of subdivisions and wavelets. The method also enabled the localising of high detail to the areas requiring highest resolution by selective face partitioning. Some visual examples of demonstrating these wavelets in action are contained in [120, 125]. Swelden's second generation wavelets were based on similar principals to the ideas of Lounsbery. However, rather than applying wavelets to manipulate the shape of their domain, the wavelets approximated functions over a fixed surface. For instance, Swelden constructed his second generation wavelets over triangular partitioning of the sphere. These were successfully applied in applications ranging from representing geophysical data [126] to resolving reflection functions and illumination on the sphere [127].

In this chapter two new spherical wavelet bases are developed based on the principals of second generation wavelets and are used to discretise the angular variable of the Boltzmann transport equation. The wavelet construction follows closely to Swelden's second generation spherical wavelets. The bases approximate with linear and quadratic variation over a triangular mesh of the octahedron's surface.

There are many potential benefits in developing and applying wavelet bases to discretise the angular variable of the transport equation. Not only have they been demonstrated to efficiently represent general spherical functions but they also possess all the necessary ingredients for incorporating angular adaption for this specific application. That is, the wavelet approximation is hierarchical and the expansion's functions have compact sup-

port. Compactly supported functions enables one to select only the wavelets required for an accurate angular representation for each spatial node of the spatial discretisation, i.e. if the flux information is concentrated in specific directions then only the wavelets with supports in these directions are required. The wavelets that do not contribute to the construction of the angular flux can simply have their expansion coefficients set to zero and removed from the calculation. This results in a reduced angular expansion which significantly increases computational efficiency. Although it may seem that the angular flux is represented by a different wavelet set at each spatial point, the hierarchy makes all angular expansions theoretically equivalent. Therefore, one has adaptivity without having the problem of a spatially varying angular approximation. A spatially varying angular approximation would presented an unfortunate situation posing various problems. This is also a reason why adaptive finite element methods on the sphere have not established themselves as a method of discretising the angular variable of the transport equation. In addition to this, the wavelet bases could also have an advantage over the existing S_N and P_N methods. The discretisations should show a significant reduction of ray-effect oscillations in the scalar flux solutions in comparison to S_N. This is because the wavelets bases use higher order approximating functions which incorporate coupling through the discretised streaming term. Furthermore, as these bases are limited to linear and quadratic representations, they should be less susceptible than P_N to Gibbs oscillations over areas of large variation in the angular flux on the unit sphere. In essence, the wavelet solutions offer a happy compromise between the methods of lowest order S_N (with ray-effects) and spectrally accurate P_N (with Gibbs oscillations).

The sections of this chapter are set out as follows. In section 3.2 a brief recap of the angular discretised Boltzmann transport equation is given. In section 3.3 the linear and quadratic octahedral wavelets are developed. Section 3.4 demonstrates the wavelets capabilities by solving a number of 1-D and 2-D steady-state neutron transport problems. Section 3.5 finishes this chapter with a conclusion.

3.2 The First Order and Angularly Discretised Transport Equation

This chapter begins with the first order Boltzmann transport equation that is discretised in angle,

$$\frac{1}{v}A\frac{\partial\Psi(r,t)}{\partial t}+A_x\frac{\partial\Psi(r,t)}{\partial x}+A_y\frac{\partial\Psi(r,t)}{\partial y}+A_z\frac{\partial\Psi(r,t)}{\partial z}+\mathcal{H}\Psi(r,t)=\mathcal{S}(r,t), \quad (3.2)$$

using an arbitrary angular approximation scheme,

$$\psi(r,\hat{\Omega},t)\approx\sum_{j=1}^{\mathcal{M}}\Psi_j(r,t)\mathcal{G}_j(\hat{\Omega}). \quad (3.3)$$

The term $\Psi(r,t)$ denotes the vector of angular coefficients corresponding to expansion (3.3). The term $\mathcal{S}(r,t)$ denotes the angular discretised source and A, A_x, A_y, A_z and $\mathcal{H}$ denote the angular mass, angular Jacobians and scattering/removal matrices. All these terms are fully discussed and defined in section 2.2.

3.3 Linear and Quadratic Spherical Octahedral Wavelets

This section develops the linear and quadratic spherical wavelets bases. For both wavelet bases, it has been decided that the approximations are to be constructed on the surface of the unit octahedron shown in figure 3.1. The advantage in taking this approach is that the flat triangular surfaces constructing the octahedron allow a method for developing the wavelet functions using standard triangular finite element techniques. Using the octahedral surface also serves to reduce the complication of discretising the angular domain in order for the wavelets to be defined. The use of the octahedral surface to represent the discretisation of particle travel, as opposed to the unit sphere, does require a mapping from the sphere to the octahedron. Therefore, a simple projection mapping is also developed here in order that the wavelet bases can be used to represent the angular flux.

The following sections are outlined as follows. Section 3.3.1 presents the projection mapping from the octahedral domain to the surface of the unit sphere. In Section 3.3.2 a

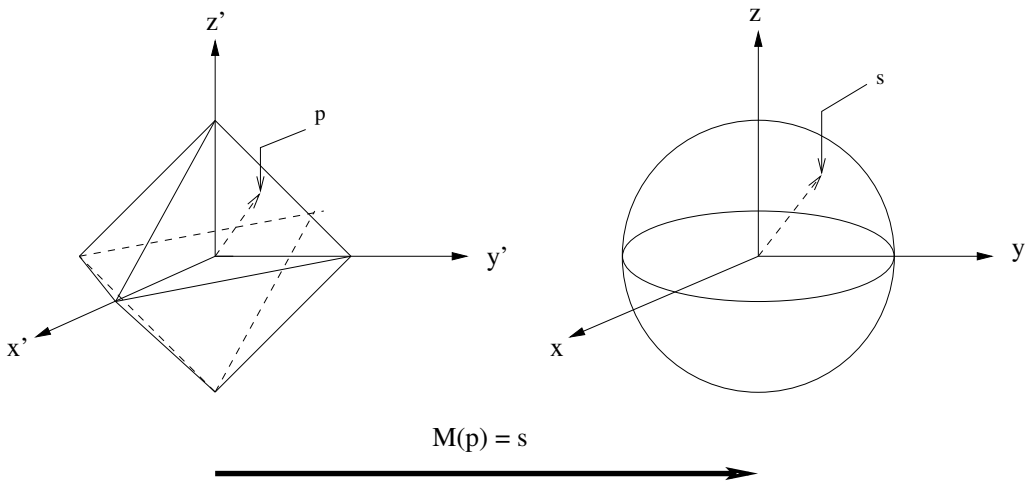

Figure 3.1: The visualisation of the normalising mapping from the octahedron to the sphere.

subdivision scheme (or triangulation) on the faces of the octahedron is presented. This scheme is required for the development of a Multiresolution Analysis (MRA) which is defined and constructed in section 3.3.3. Section 3.3.4 constructs the linear and quadratic wavelets relating to MRA developed in section 3.3.3. In Section 3.3.5, a modification to the wavelet functions is presented in order to mitigate ray effect-like errors destroying 1 and 2 dimensional solutions which were first reported by Buchan in [92, 93].

3.3.1 Representing the Angular Flux over the Octahedron

Representing the direction of particle travel on the octahedron's surface requires a mapping from the sphere S^2 to the octahedron's surface O. The mapping used in this chapter is based on the compression of the sphere over the polyhedron through the normalising projection M. That is, the projection $M : O \rightarrow S^2$ is defined by $M(p) = s = p/\|p\|$, where $s \in S^2$ and $p \in O$. A full illustration of this mapping is presented in figure 3.1. The mapping is an isomorphism that guarantees the existence of a function $F : O \rightarrow R$, that is equivalent to any spherical function $f : S^2 \rightarrow R$ through the relationship $F(p) = f(M(p)) = f(s)$.

In order to proceed a number of alterations to the angular discretised transport equation (3.2) are required . First all spherical functions, including the trial functions G_i of (3.3), are to be defined on the octahedron. Second, the definitions of the angular matrices presented in (2.22) and (2.39) are to be integrated over the octahedron's surface. This is defined by,

$$\int d\mu \int d\omega \; f(\mu,\omega) = \int dp \int dq \; f(\mu(p,q),\omega(p,q)) \; |J(p,q)| \tag{3.4}$$

$$= \int dp \int dq \; F(p,q) \; |J(p,q)|, \tag{3.5}$$

where f represents an arbitrary spherical function. The variables p and q denote two position variables over the octahedron and the term $|J(p,q)|$ represents the Jacobian of the mapping from the spherical coordinates (μ,ω) to the coordinates of the octahedron (p,q), i.e. the determinant of the Jacobian matrix between the two spaces.

$$|J(p,q)| = \begin{vmatrix} \frac{d\mu}{dp} & \frac{d\omega}{dp} \\ \frac{d\mu}{dq} & \frac{d\omega}{dq} \end{vmatrix}$$

As an example, applying (3.4) to the integrals (2.22) re-defines the angular Jacobian matrix A_x as,

$$A_{x\;ij} = \int dp \int dq \; \tilde{\Omega}_x(p,q) \; G_i(p,q) \; G_j(p,q) \; |J(p,q)|,$$

where the angular basis functions G_i and G_j are functions on O and $\tilde{\Omega}_x$ satisfies,

$$\tilde{\Omega}_x(p,q) = \Omega_x(M(p,q)) = \Omega_x(\mu,\omega).$$

The complexity of this mapping will require a numerical scheme for calculating the angular integrals. This is covered in a later section where a method is presented that gives a high order and consistently accurate numerical scheme for calculating the integrals when a wavelet expansion is used.

3.3.2 Subdivision Scheme on the Octahedron

Construction of second generation wavelets begins with a recursive partitioning of the domain on which the wavelets are defined. A subdivision scheme defined on the octahedron

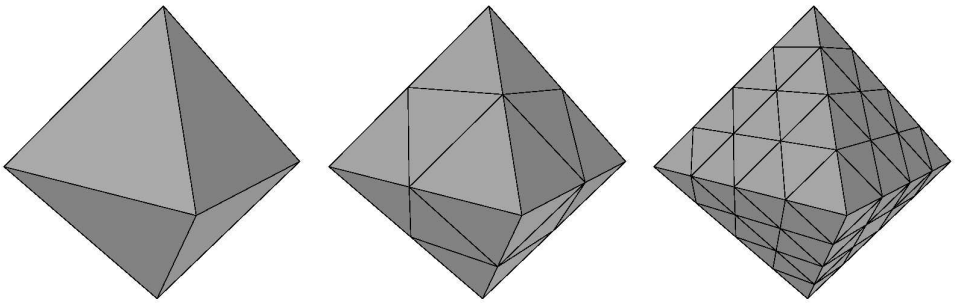

Figure 3.2: The octahedron (left) with the first (middle) and second (right) subdivisions.

is used to recursively divide every face into four new identical faces. This is illustrated in figure 3.2 for which the faces of the initial octahedron have undergone one and two subdivisions. The faces generated by i subdivisions are denoted by O_i, with O_0 denoting the 8 original faces.

In the theory that follows the treatment for the linear and quadratic wavelets will vary from case to case. However, the same notation indexing points has been used for both wavelets and so the reader should be aware of the differences in indexing between the methods.

For the linear wavelets, the set $S_j, j \in \{0, 1, \ldots, \infty\}$, denotes the vertices generated by j subdivisions on the octahedron. For the quadratic wavelets, the set $S_j, j \in \{0, 1, \ldots, \infty\}$, denotes the vertices and mid-edge points upon j partitions on the octahedron. Figure 3.3 presents an illustration of the sets S_j for both linear and quadratic cases. For both cases, the sets S_j are formally defined as,

$$S_j = \{s_{j,k} \in O, k \in K(j)\}, \qquad j \in \{0, 1, \ldots, \infty\},$$

where $K(j)$ is an index set relating to the vertices (linear case) or the vertices and mid-edge points (quadratic case) of the j^{th} partition of the octahedron. As the sets of S_j are nested, $S_j \subset S_{j+1}$, the index set $K(j)$ is ordered so that it is also nested. Therefore, at any point $s_{j,k} \in S_j,\ k \in K(j)$, the following relationship is made to hold, $s_{j,k} = s_{j+1,k},\ k \in K(j+1)$.

Finally, the set $M(j)$ is used to index the new elements of S_{j+1} created at the $j+1$

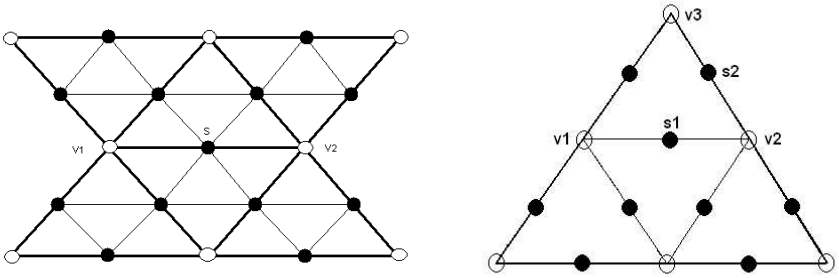

Figure 3.3: These diagrams show part of the octahedron's face partitioned to levels j and $j+1$ - indicated by the bold and thin lines respectively - for the linear (left figure) and quadratic (right figure) wavelets. The hollow dots represent the points of S_j, these are indexed by the set $K(j)$. The solid dots represent the new points of S_{j+1}, these are indexed by the set $M(j)$. For both diagrams, the point $s1$ indicates the position of a lazy wavelet of level j and $v1$ and $v2$ indicate the positions of the two scaling functions used in its lifting. For the quadratic wavelet, the lazy wavelet associated with point $s2$ will only have the scaling function associated with the point $v2$ used in its lifting. The scaling function associated with the point $v3$ has a zero integral.

subdivision. That is, $M(j) = K(j+1) \setminus K(j)$, where '\' is the set removal operator.

3.3.3 Multiresolution Analysis on the Octahedron

A multiresolution Analysis (MRA) is a nested sequence of subspaces of the complete set of Lebesgue integrable functions $L_2(X)$, for some space X. In the context of functions on the surface of the octahedron, the space $L_2(O)$ is defined as the set of functions such that for every $f \in L_2(O)$, the condition $\int_O f^2\, do < \infty$ holds, where do is the measure of area over the octahedron. An MRA based on the octahedron is defined as a nested sequence of closed subspaces $V_j \subset L_2(O),\ j \in \{0, 1, \ldots, \infty\}$, that satisfy the following conditions:

C1 $V_0 \subset \ldots V_{j-1} \subset V_j \subset V_{j+1} \ldots \subset L_2(O)$

C2 $\bigcup_{j=0}^{\infty} V_j$ is dense in $L_2(O)$

C3 At each level j there exists a set of scaling functions $\varphi_{j,k}, k \in K^{'}(j)$ for some index set $K^{'}$, for which the set $\{\varphi_{j,k},\ \ k \in K^{'}(j)\}$ forms a Riesz basis of V_j.

In developing the linear and quadratic wavelets, two separate MRAs are required. Their construction, together with the set of scaling functions of condition C3, are defined in turn before returning to the general properties of a multiresolution analysis and their application in numerical analysis.

Linear MRA

For the linear MRA, the spaces V_j are defined as the collection of all functions which are piece-wise linear over the faces generated on j subdivisions of the octahedron. The spaces V_j satisfy the conditions of the MRA as the subdivision scheme ensures they are nested (since the faces generated on j subdivisions can be created from the faces generated on $j+1$ subdivisions) and that V_j is dense on the space $L_2(O)$ (since in the limit of j the sub-divided faces become arbitrarily small and so there exists a function in V_j, for some

j, that is arbitrarily close to any function in $L_2(O)$). In order to complete the MRA, the set of scaling functions $\varphi_{j,k}$, that span V_j, are defined as the linear triangular finite element trial functions [34]. These trial functions are associated with the nodal points in the triangular mesh of O_j. They have the value of one over their associated node and attenuate linearly towards their immediate neighbouring vertices. Since the scaling functions $\varphi_{j,k}$ are associated with the nodal points generated by j partitioning of the octahedron, the index set $K^{'}(j)$ in C3 is identical to the index set $K(j)$ defined by the subdivision scheme (linear case). The set of scaling functions for the linear MRA can be specifically defined as,

C3 At each level j there exists a set of scaling functions $\varphi_{j,k}, k \in K(j)$, for which the set $\{\varphi_{j,k}, \ \ k \in K(j)\}$ forms a Riesz basis of V_j. The set $K(j)$ is the index set generated by the subdivision scheme on the octahedron (linear case) and the scaling functions are defined as the standard linear Finite Element trial functions associated with the points indexed by $K(j)$.

Quadratic MRA

For the quadratic MRA, the spaces V_j are defined as the collection of all functions that are piece-wise quadratic over the faces generated by j subdivisions of the octahedron. Using the same arguments used in the linear case, the spaces V_j can be shown to satisfy the first two conditions of the MRA. For the third condition, the set of triangular quadratic finite element trial functions [34] on O_j are employed as the scaling functions $\varphi_{j,k}$ that generate a basis for the spaces V_j. The scaling functions are associated with the vertices and mid-edge positions of the triangular mesh generated by O_j. They have the value 1 over their respective collocation point and are 0 over the other nodal and mid-edge positions, as illustrated in figure 3.4. Similar to the linear MRA, a relationship between the scaling functions $\varphi_{j,k}$ and the index set $K(j)$, generated from the subdivision scheme (quadratic case), can be established. That is the scaling functions $\varphi_{j,k}$ are associated with the points indexed by the elements of $K(j)$. The scaling functions of the quadratic MRA can therefore be specifically defined by C3 (linear case) with the adjustment of referring to the quadratic set $K(j)$ and quadratic finite element trial functions.

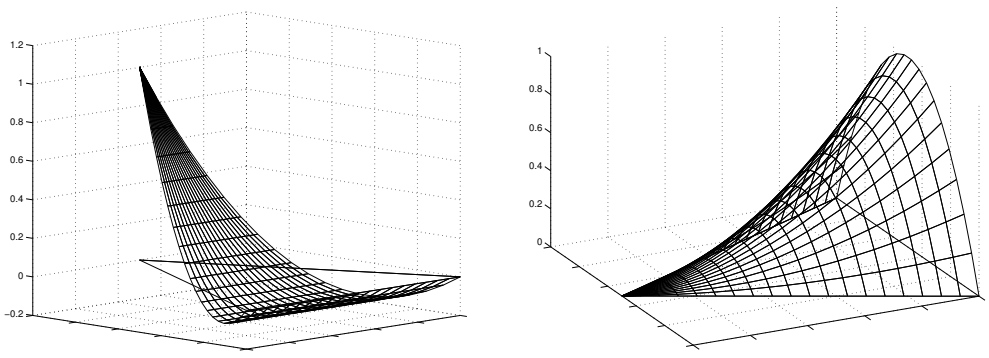

Figure 3.4: The quadratic finite element trial functions over a triangular element. Left: Corner function, Right: mid-edge function

General Properties of the MRA and their application in Numerical Analysis

This section lists some important general properties of an arbitrary MRA and their associated scaling functions.

Refinement Relation:

The nested spaces V_j imply the scaling functions $\varphi_{j,k} \in V_j \subset V_{j+1}$ can be constructed from a combination of the scaling functions spanning V_{j+1}. This is referred as the refinement relation,

$$\varphi_{j,k} = \sum_{m \in K(j+1)} h_{j,k,m}\, \varphi_{j+1,m}. \tag{3.6}$$

For both linear and quadratic MRAs, the interpolating scaling functions, that is $\varphi_{j,k}(s_{j,l}) = \delta_{k,l},\ \ \forall j \in \{0, 1, \ldots\ \infty\},\ \forall\, k, l \in K(j)$, lead to the simplification of the refinement relation,

$$\varphi_{j,k} = \varphi_{j+1,k} + \sum_{m \in M(j)} h_{j,k,m}\varphi_{j+1,m}, \tag{3.7}$$

where the right hand side summation involves only the scaling functions associated with the collocation points indexed by $M(j)$. Furthermore, from the scaling function's interpolating property, the coefficients $h_{j,k,m}$ are calculated to be the value of $\varphi_{j,k}$ at the point $s_{j+1,m}$, that is, $h_{j,k,m} = \varphi_{j,k}(s_{j+1,m})$.

Numerical applications:

The purpose of the MRA's nested spaces V_j are to provide spaces for the projection of an approximating of a function $f \in L_2(O)$. The projected function $f_j \in V_j$, for any $j \in \{0, 1, \dots \infty\}$ is represented by an expansion of the scaling functions $\varphi_{j,k}$ that span V_j,

$$f \to f_j = \sum_k \lambda_k \, \varphi_{j,k}, \tag{3.8}$$

where the values λ_k denote the expansion coefficients.

The accuracy of the approximation is governed by the selection of j and property C2 indicates that f is fully recovered by letting $j \to \infty$. This dense conditioning of the spaces V_j (as $j \to \infty$) creates a basis for $L_2(O)$.

$$L_2(O) = \lim_{j \to \infty} V_j \tag{3.9}$$

3.3.4 Wavelets Generated by the Lifting Scheme

The purpose of wavelets are that they create bases for the spaces W_j that complements V_j in V_{j+1}. That is, $V_{j+1} = V_j \oplus W_j$ for all $j \in \{0, 1, \dots \infty\}$. The set of wavelet functions that span W_j are denoted as,

$$\{\psi_{j,m}, \;\; m \in M'(j)\}, \tag{3.10}$$

where $M'(j)$ denotes some index set.

The complementing condition of W_j and V_j imply the wavelet functions $\psi_{j,m}$ are contained in the space V_{j+1}, i.e. $\psi_{j,m} \in W_j \subset V_{j+1}$. This property implies a similar relationship to the refinement relation (3.6) exists for the wavelet functions. Each wavelet $\psi_{j,m}$ can therefore be expanded in terms of the scaling function spanning V_{j+1},

$$\psi_{j,m} = \sum_{k \in K(j+1)} v_{j,m,k} \, \varphi_{j+1,k}.$$

By recursively applying the relationship $V_j = V_{j-1} \oplus W_{j-1}$, the space V_j can be decom-

posed into the spaces V_0 and W_k for $k \in \{0, 1, \ldots, j-1\}$.

$$\begin{aligned} V_j &= V_{j-1} \oplus W_{j-1} \\ &= V_{j-2} \oplus W_{j-2} \oplus Wj-1 \\ &= \ldots \\ &= V_0 \ \oplus_{k=0}^{j-1} W_k \end{aligned} \tag{3.11}$$

Furthermore, combining (3.9) and (3.11) decomposes the Lebesgue space $L_2(O)$ into the spaces W_j, for $j \in \{0, 1, \ldots, \infty\}$, and V_0.

$$L_2 = V_0 \oplus_{j=0}^{\infty} W_j$$

This creates a second basis spanning $L_2(O)$ using all wavelet bases for W_j and the lowest order scaling functions spanning V_0. This is explicitly written as,

$$\{\varphi_{o,k}, \ \ k \in K(0)\} \cup \{\psi_{j,m}, \ \ j \geq 0, m \in M(j)\}.$$

Therefore, any function $f \in L_2(O)$ can be expanded by,

$$f = \sum_{k \in K(0)} \lambda_{0,k} \, \varphi_{0,k} + \sum_{l=0}^{\infty} \sum_{k \in M(l)} \gamma_{l,k} \, \psi_{l,k}, \tag{3.12}$$

and, similarly, the projected approximation $f_j \approx f$ in (3.8) can also be represented by,

$$f_j = \sum_{k \in K(0)} \lambda_{0,k} \, \varphi_{0,k} + \sum_{l=0}^{j-1} \sum_{k \in M(l)} \gamma_{l,k} \, \psi_{l,k}. \tag{3.13}$$

Linear and Quadratic Octahedral Wavelets

A simple approach in constructing a basis for the complementing space W_j is to use,

$$\psi_{j,m} = \varphi_{j+1,m}, \qquad \forall \, m \in M(j). \tag{3.14}$$

This defines the basis for the space W_j to be the scaling functions of the space V_{j+1} that are associated with the new collocation points indexed by the set $M(j)$. The wavelets constructed using (3.14) were coined 'lazy wavelets' by Swelden since they required no extra computation for their construction. However, it has also been shown that this approach does not guarantee a stable basis for W_j [126].

The wavelets are therefore re-defined using the modification,

$$\psi_{j,m} = \varphi_{j+1,m} - \sum_{k \in K(j)} S_{j,m,k}\, \varphi_{j,k} \qquad m \in M(j), \tag{3.15}$$

where contributions from the scaling functions $\varphi_{j,k}$ are removed from (3.14), a process which is a generalisation of the lifting scheme [90, 119]. Arbitrary values for $S_{j,m,k}$ can be used in (3.15) since the complementing properties between W_j and V_j always hold, and so they are selected in order to generate more stable bases. The approach used here is to ensure that all wavelets have vanishing integrals (that is, integration of each wavelet over the octahedron equates to zero) as this has been show to increase their stability [90, 119]. In addition to this, it is important to keep the support of the wavelets small in order that the functions remain localised. This requires restricting the lifting contribution from the scaling functions in (3.15), to only the two functions $\varphi_{j,k1}$ and $\varphi_{j,k2}$, where $k1, k2 \in K(j)$ index the adjacent vertices of $m \in M(j)$. This is illustrated in figure 3.3.

For the linear wavelets, the values of $S_{j,m,k}$ are defined as,

$$S_{j,k,m} = \begin{cases} \frac{\int \varphi_{j+1,m}}{2 \int \varphi_{j,k}} \text{ if } S_{j,k} \text{ is a neighbour of } S_{j+1,m} \\ 0 \text{ otherwise} \end{cases}$$

which generate functions of the form presented in figure 3.5.

For the quadratic wavelets, the above definition is used for the values $S_{j,m,k}$ when both the scaling functions used in lifting have a non zero integral. In the case where an adjacent scaling function does have a zero integral (see figure 3.3), the second scaling function performs the full lifting. In this situation the set $S_{j,m,k}$ is defined by,

$$S_{j,k,m} = \begin{cases} \frac{\int \varphi_{j+1,m}}{\int \varphi_{j,k}} \text{ if } S_{j,k} \text{ is the only neighbour of } S_{j+1,m} \text{ with } \int \varphi_{j,k} > 0 \\ 0 \text{ otherwise.} \end{cases}$$

3.3.5 A Modified Wavelets Basis

The wavelet described in the previous sections have been shown to be highly inaccurate when resolving certain 1-D neutron transport problems [92, 93]. These inaccuracies were essentially a ray-effect type error that caused flux to travel in beams perpendicular to the

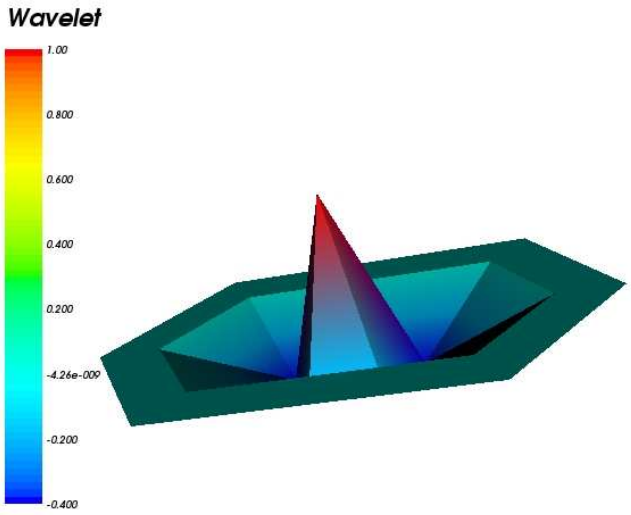

Figure 3.5: A diagram of a typical lifted linear wavelet on a sub-section of the surface of the octahedron.

spatial dimension of the 1-D problem. As shown in the first numerical example of this chapter this results in a build up of flux over source regions and discontinuous jumps in the solutions appear at the source edges. The investigation of [92, 93] showed that rotating the octahedron could reduce the inaccuracies to some extent. However, this method could only be used as a temporary measure since rotation was not effective for high order wavelet expansions.

In order to completely mitigate these effects the wavelet expansions must allow discontinuities in its approximation at the edges of the original octahedron (which lie in direct line with the x, y and z axes) as this offsets the ray-effects from directing flux in the angles that cause the errors to occur. It is simple to enforce this discontinuity by defining the scaling functions of the MRA space to be the finite element functions associated with the elements of the subdivided octahedron (as before). However, for a node that lies along the edges of the original octahedron O_0, we group its adjacent elements into sets, so that the elements of each set belong to the same face of the original octahedron. Then, rather than using a single function to interpolate over all adjacent elements of the node, separate

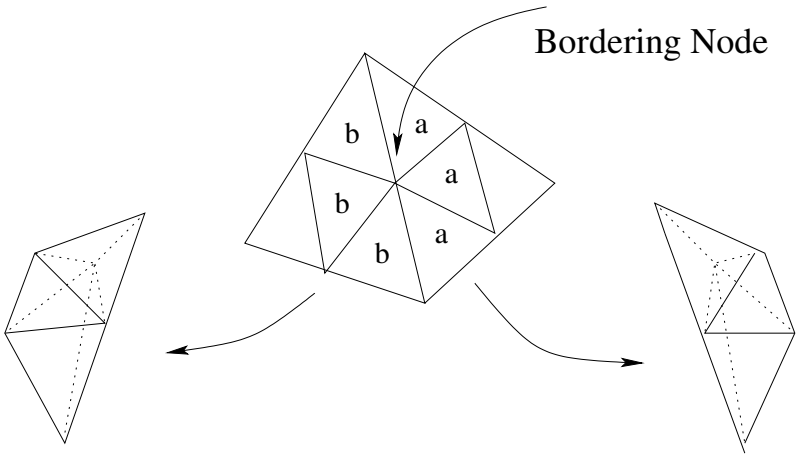

Figure 3.6: This diagram illustrates how the modified wavelet bases are generated using the modification developed in section 3.3.5. The diagram shows two faces of the octahedron that has been subdivided once. The vertex being considered is positioned on the edge of the two octahedral faces. In order to generate a discontinuity between the two octahedral faces the subdivided elements are partitioned into two sets, denoted a and b. Each set groups together the elements that belong to the same face of the octahedron. Two septate finite element functions associated with the node are then used to span the two element sets - this illustrated by the separate functions indicated by the two arrows (the dotted lines show the profile of the linear finite element functions). These functions, together with the continuous finite element functions on the internal vertices of the octahedron, will generate the desired approximating spaces of the MRA.

finite element functions are used to interpolate over each subset of elements. That is, the new functions represent the contribution of the original finite element function over the separate faces of the octahedron. This allows discontinuity in the approximation at the edges of the octahedron - see figure 3.6 for an illustration. The wavelet bases can then be generated using the derivation given above.

In order to make these amendments clear, and to avoid confusion in the numerical examples, our DEFAULT wavelet expansions used are those which give continuous approximations across the elements of the subdivided octahedron, except when an edge between two elements form part of the edge of the original octahedron O_0. In which case discontinuity

is allowed. In the first numerical example these default wavelets are compared with the original wavelets that gave approximations which were continuous over the whole octahedron. This was so that the modifications developed in this sections could be demonstrated to be work properly. In addition to this amendment of allowing discontinuity over the edges of the octahedron we also investigate (in the second numerical example) a second wavelet expansion that allows discontinuity across all elements of the subdivided octahedron. That is, discontinuous finite element functions are used to span the spaces of the MRA (and again the wavelet are developed in the same way). In order to distinguish between the two wavelet expansions in this example, the default wavelets are said to be continuous (as they are continuous everywhere but for the edges of the octahedron) and the other wavelet expansion is said to be discontinuous.

However, the reader should note that unless otherwise stated, the wavelets being considered will be the default wavelets.

3.3.6 Calculating the Angular Integrals

In order to accurately approximate the angular integrals defined by equations (2.22) and (2.39), the following method is advised when using the spherical wavelet expansions.

Initially, the angular matrices of equation (2.21) are calculated in terms of the highest order scaling functions, which are essentially a finite element expansion over the elements of the subdivided octahedron. The integrals over the scaling function space are approximated using a high order Gauss quadrature scheme. This is required to ensure the information of the angular variation (that is the $\hat{\Omega}_x, \hat{\Omega}_y$ and $\hat{\Omega}_z$ terms) are accurately captured. To integrate on the triangular elements, it is advisable to map the elements onto a unit quadrilateral, this is discussed in [128] where two nodes of the quadrilateral are collapsed to make the triangle, and to use two Gauss quadratures to resolve the two dimensions of the quadrilateral element.

The angular matrices can then be mapped into their corresponding wavelet space using the mapping defined in (3.15). For each level (l say) above the lowest order approximation,

two mappings are required. The first is to map the correct higher order scaling functions to the scaling functions of level $l-1$. The second mapping is required to map the remaining higher order scaling function into their lifted wavelet form. Both mappings can be defined as,

$$\psi_i^{new} = \sum_j a_i^j \psi_j^{old} \tag{3.16}$$

where ψ_i^{new} and ψ_j^{old} denote a new and old scaling/wavelet function, respectively, and a_i^j denote the mapping coefficients. In its general form the mapping can be represented for all $\mathcal{M}$ angular variables by,

$$\psi_{new} = W\psi_{old}, \tag{3.17}$$

where W is an $\mathcal{M} \times \mathcal{M}$ matrix defined defined by $W_{i,j} = a_i^j$. The angular matrices can then be mapped in order to define the angular discretised transport equation in terms of ψ^{new} by pre-multiplying and post multiplying them by W and W^T, respectively. For example the angular Jacobian A_x becomes,

$$A_x \rightarrow W A_x W^T. \tag{3.18}$$

There are two advantages in taking this approach. The first is that the numerical integration is of high order accuracy and that the order of accuracy is consistent for all wavelet functions. Note here that the $\alpha_i^{\cdot,l,m}$ are also computed to a high degree of accuracy. This in turn helps mitigate any non-conservation of scattered particles when approximations of the integrals of equations (2.39) are made. The second advantage is that the computation of the angular information is inexpensive. This is because there are relatively few integrals to perform in the high order finite element space (compared to the number of angular functions) and that the mapping matrices W will also be sparse - thus few calculation are required for the mapping (3.18).

3.4 Numerical Examples

In this section a number of steady state transport problems are solved in order to demonstrate the accuracy of the octahedral wavelets. The first numerical examples aim to demonstrate the wavelets capabilities on solving 1-D test problems. The first example

is designed specifically to demonstrate the effects of the modification procedure detailed in section 3.3.5 on solving a simple pure absorbing problem. This is followed by solving the more demanding Reed Cell Problem using the modified wavelet expansions. In this example the option of using the fully discontinuous wavelets, also described in section 3.3.5, are investigated.

This section also solves a number of 2-D steady state problems. For these demonstrations, 2 box problems, with and without isotropic scattering, a straight duct problem and a source within a void problem are solved using only the default wavelets. The aims of these numerical examples are to investigate the wavelet expansions for their accuracy and their susceptibility to ray effects forming in their solutions. The approximations are compared with the solutions from the established S_N and P_N discretisations. These numerical examples are also designed to show the wavelets working within a finite element spatial discretisation framework. For this demonstration both structured and unstructured finite element meshes are used.

All wavelet solutions are compared to solutions provided by established methods using established numerical codes. Benchmark solutions have been provided by EVENT [129] - an even parity code using a spherical harmonics angle discretisation and finite element spatial discretisation, MCNP [130]- a probabilistic Monte Carlo code and a discontinuous Galerkin Discrete Ordinate technique (DG-DO) [21].

For the following numerical examples the linear and quadratic wavelet solutions relating to N partitions of the octahedron are denoted by LW_N and QW_N, respectively. The sizes of the wavelet discretisations are listed in table 3.1.

3.4.1 1-D Numerical Examples

In the following 1 dimensional numerical examples the spatial variable is approximated using a control volume spatial discretisation. This discretisation method was described in section 2.3.3 where estimations of the angular flux at cell faces and domain boundaries are approximated using a Riemann technique [89, 87].

Angular Expansion sizes		
N (Angular expansion)	OLW_N	OQW_N
0	12	24
1	24	60
2	60	180
3	180	612
4	612	

Table 3.1: This table lists the angular expansion sizes of the octahedral linear and quadratic wavelet discretisations.

Problem 1: Isotropic Source within a Pure Absorbing Material

This problem was designed to demonstrate the capabilities of the modified wavelets detailed in section 3.3.5 (that is where the approximations allow discontinuity along the face edges of the octahedron - the default wavelets) for mitigating the ray-effect type inaccuracies described in reports [92] and [93]. The problem domain, presented in figure 3.7, is 10cm in length and contains a pure absorbing material with a cross-section of $0.5\mathrm{cm}^{-1}$. An isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ occupies the central 2cm region. Vacuum boundaries conditions are prescribed along both edges and the spatial domain was discretised using a regular mesh of 200 control volumes. The problem was solved using the modified LW_1 (the default wavelets), un-modified LW_1 and P_9 angular discretisations. Their solutions are presented in figure 3.8.

The results reveal the un-modified wavelet solution possess ray-effect type inaccuracies in the form of a discontinuous jump in the scalar flux solution over the source region. These are clearly seen to be un-physical discontinuities when the solution is compared to the P_9 approximation. However, the results also show the modified wavelets to have mitigated the errors and provided a solution that compares closely to the P_9 approximation. It is therefore evident from this numerical example that the method developed in section 3.3.5 does mitigated the ray-effect inaccuracies that form in the un-modified wavelets solutions.

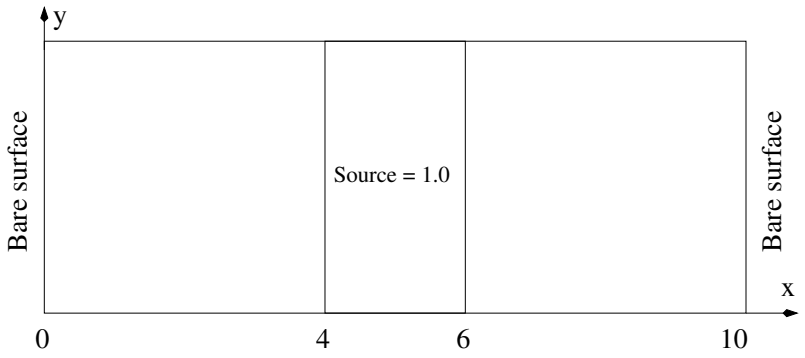

Figure 3.7: Problem 1: Diagram of the simple 1-D problem. The setup has a constant pure absorbing cross-section of $0.5\mathrm{cm}^{-1}$ across the whole domain, vacuum surfaces bounding both sides and an isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ occupying the central 2cm region.

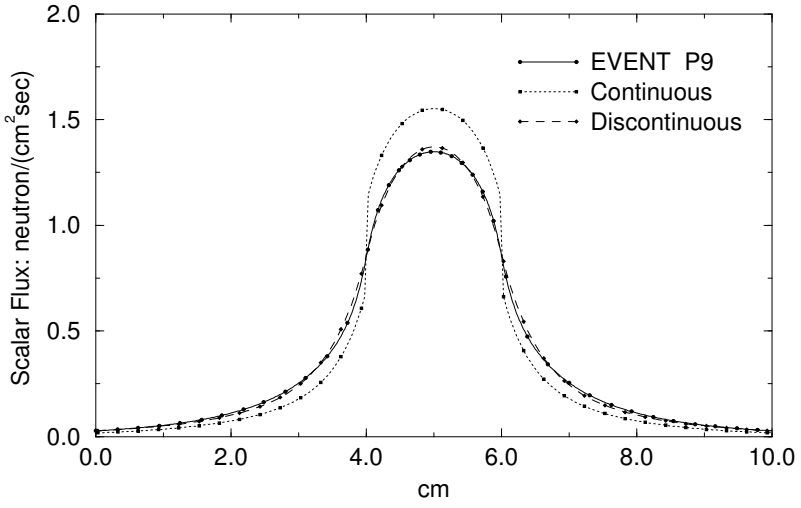

Figure 3.8: Scalar flux solutions obtained from the linear wavelets with and without the discontinuity over the octahedron's faces. Calculations were performed on problem 1.

Figure 3.9: Problem 2: Diagram of the 1-D Reed Cell Problem.

Problem 2: The Reed Cell Problem

This problem is designed to demonstrate the capabilities of the wavelets in approximating demanding 1-D geometries. In this demonstration a comparison between the default linear and quadratic wavelets solutions is given. The final part of this section gives a demonstration of the fully discontinuous linear wavelet's which were discussed in section 3.3.5.

For this demonstration the Reed cell problem is solved. The problem domain is presented in figure 3.9 and the corresponding source and material information is listed in table 3.2. The problem comprises of 5 regions that lies within a reactor lattice cell that is located at the edge of a bare core. Region 1, positioned next to a reflecting boundary, is 2 cm in length and contains an optically thick, pure absorbing material with a cross-section $\sigma_t = 50.0\ \mathrm{cm}^{-1}$ and a high isotropic source of intensity 50.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Next to this is a 1cm fuel can containing a pure absorbing material with a cross-section $\sigma_t = 5.0\ \mathrm{cm}^{-1}$, which is designed to cause a sharp attenuation of particles. The next region is a 2cm void, and this is followed a 3cm moderator. The moderator's material has isotropic scattering $\sigma_s = 0.9\mathrm{cm}^{-1}$ and a small absorber $\sigma_a = 0.1\mathrm{cm}^{-1}$. Region 4, the left side of the moderator, contains an isotropic neutron source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The problem is bounded at the right with a bare surface and the spatial domain discretised using a regular mesh of 800 control volumes.

Figure 3.10 presents the LW_2 and P_9 scalar flux solutions, which use 120 and 100 angular basis functions respectively. Both results are in close agreement over all 5 regions of the

Region	Source $(cm^{-2}s^{-1})$	$\sigma_a\ (cm^{-1})$	$\sigma_s\ (cm^{-1})$
1	50.0	50.0	0.0
2	0.0	5.0	0.0
3	0.0	0.0	0.0
4	1.0	0.1	0.9
5	0.0	0.1	0.9

Table 3.2: Source and material properties of Problem 2, the Reed cell problem.

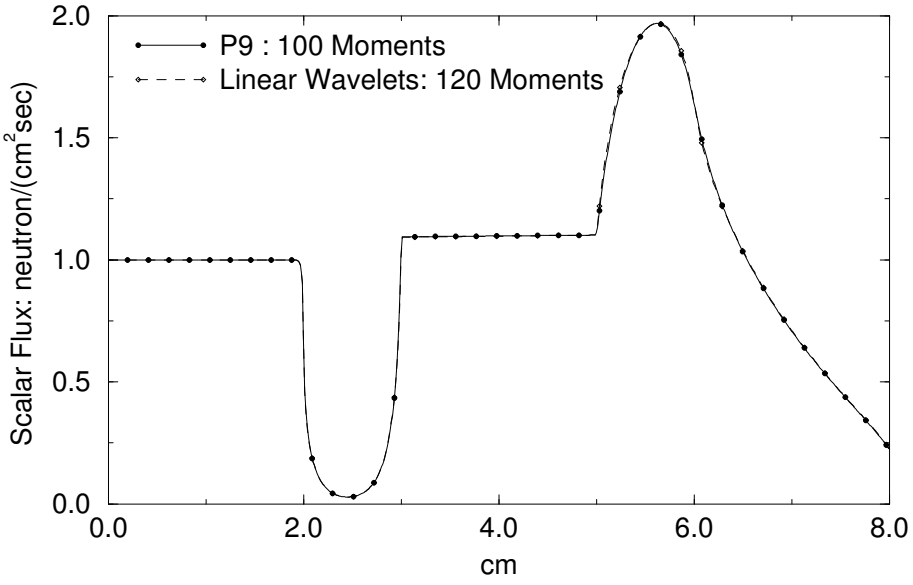

Figure 3.10: The LW_1 and P_9 scalar flux solutions of the Reed Problem.

problem domain. In order to demonstrate the solution's convergence with respect to the angular expansion, figure 3.11 presents the scalar flux solutions for the LW_1, LW_2, P_5 and P_9 angular approximations. These use 48, 120, 36 and 100 angular basis functions, respectively. The graph is concentrated over the peak in region 4, as this region shows the largest variation in the solutions, and the results are compared to the benchmark solution provided by EVENT using P_{19}. The corresponding difference plots to these solutions are show in figure 3.12.

The LW_1 and P_5 solutions show large differences in their scalar flux profile over region 4.

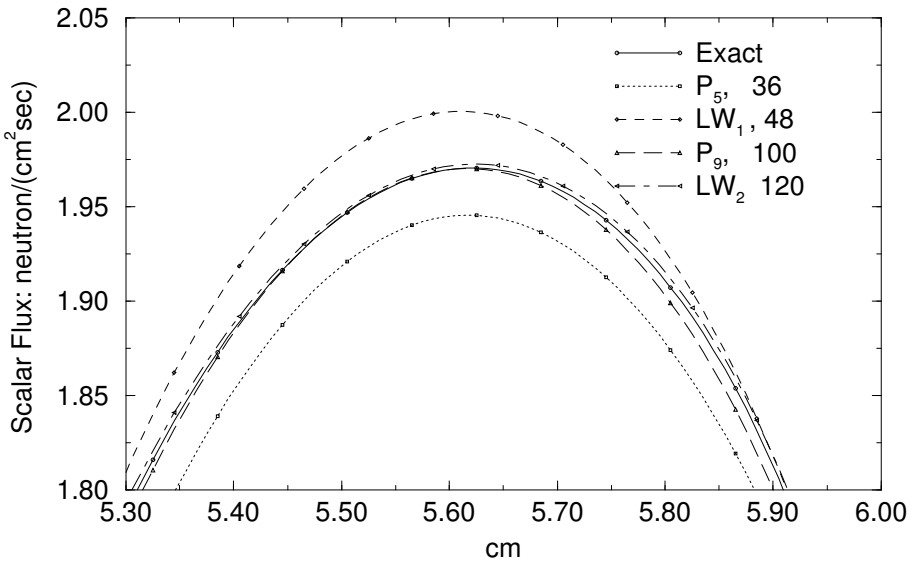

Figure 3.11: The LW_1, LW_2, P_5 and P_9 scalar flux solutions of the Reed cell problem. The graphs are concentrated over the peak of the angular flux in the moderator region.

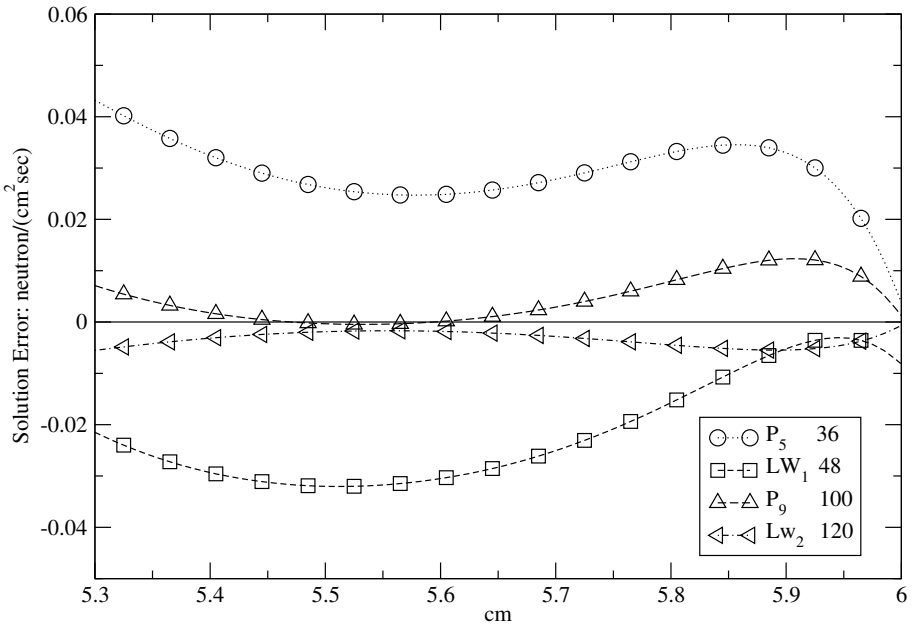

Figure 3.12: The LW_1, LW_2, P_5 and P_9 scalar flux errors for the Reed cell problem. The graphs are concentrated over the peak of the angular flux in the moderator region.

The flux peak in the LW_1 approximation reaches approximately 1.99 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, which is significantly higher than the P_5 solution that rises to only 1.94 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The benchmarked solution show the correct peak to reach 1.97 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, showing LW_1 and P_5 to be roughly equal in accuracy - this is confirmed by the difference plots in figure 3.12. The scalar flux profiles from the higher ordered LW_2 and P_9 approximations are in close agreement with the exact solution. The difference plots also show that the errors of these two solutions are roughly of equal magnitude and that convergence is close to being reached.

Figure 3.13 presents the LW_2 angular flux plots at the centre of each of the 5 regions of the Reed cell problem. These plots clearly demonstrate the wavelets capabilities in resolving a diverse range of flux distributions from isotropic fluxes in region 1 to highly directional distributions in region 2. These solutions also show the wavelets are capable resolving scattering cross-sections. This is demonstrated in region 5 where particles are seen to be propagating in the negative x direction - this can only arise through scattering. The wavelets have also performed remarkably well in resolving complex flux distributions. This is demonstrated in region 3 where the angular flux develops a ringed peak about the negative x axis. Figure 3.14 presents the first four wavelet solutions at this mid point of region 3. These illustrations demonstrate the wavelets power as each incrementation in the angular expansion adds a finer layer of detail to the solution. The peaked ring becomes clearly visible and well defined using LW_2, indicating that convergence in the solution has nearly been reached using this level of angular resolution.

Figure 3.15 presents the QW_1 and P_9 scalar flux solutions, which require 120 and 100 angular basis functions, respectively. The two results are again shown to be in close agreement over all 5 regions of the problem domain. The graph presented in figure 3.16 shows the QW_0, QW_1, P_5 and P_9 solutions that again concentrate on the peak over region 4. The results show QW_0 performs better than P_5 since the peaks reach 1.98 and 1.94 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, respectively. This is highlighted by the error plots shown in figure 3.17. This graph clearly shows the quadratic wavelet approximation to be closer to the exact solution over the region considered. The difference plots also shows QW_1 to be more accurate than P_9 and that convergence with the Benchmarked solution has almost been reached.

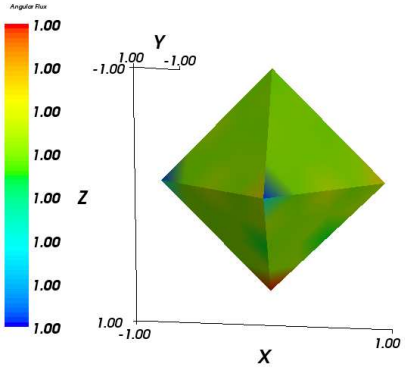

(a) Angular flux: region 1

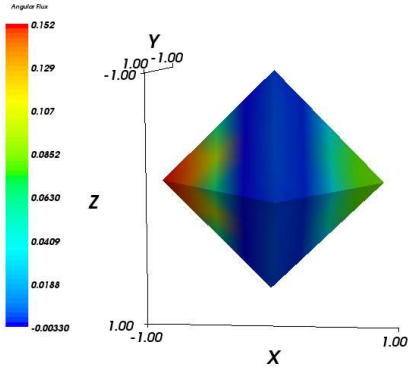

(b) Angular flux: region 2

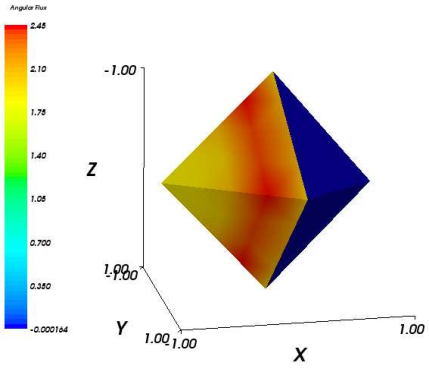

(c) Angular flux: region 3

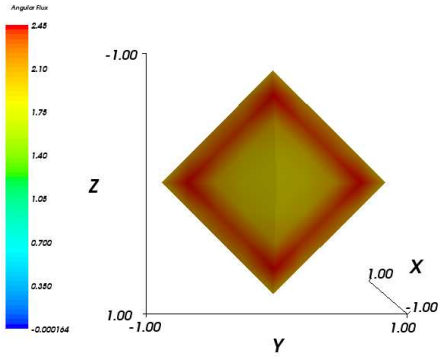

(d) Angular flux: region 3

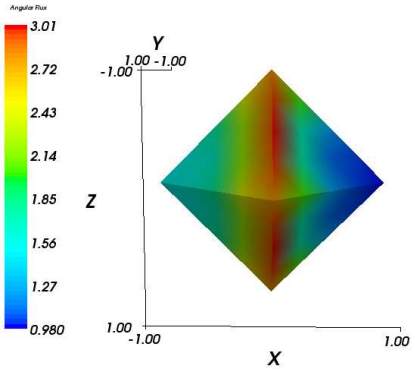

(e) Angular flux: region 4

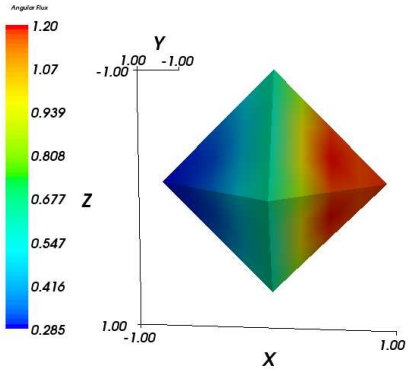

(f) Angular flux: region 5

Figure 3.13: Angular flux plots at the mid point of the 5 regions of the Reed cell problem. The fluxes were generated by the LW_2 discretisation.

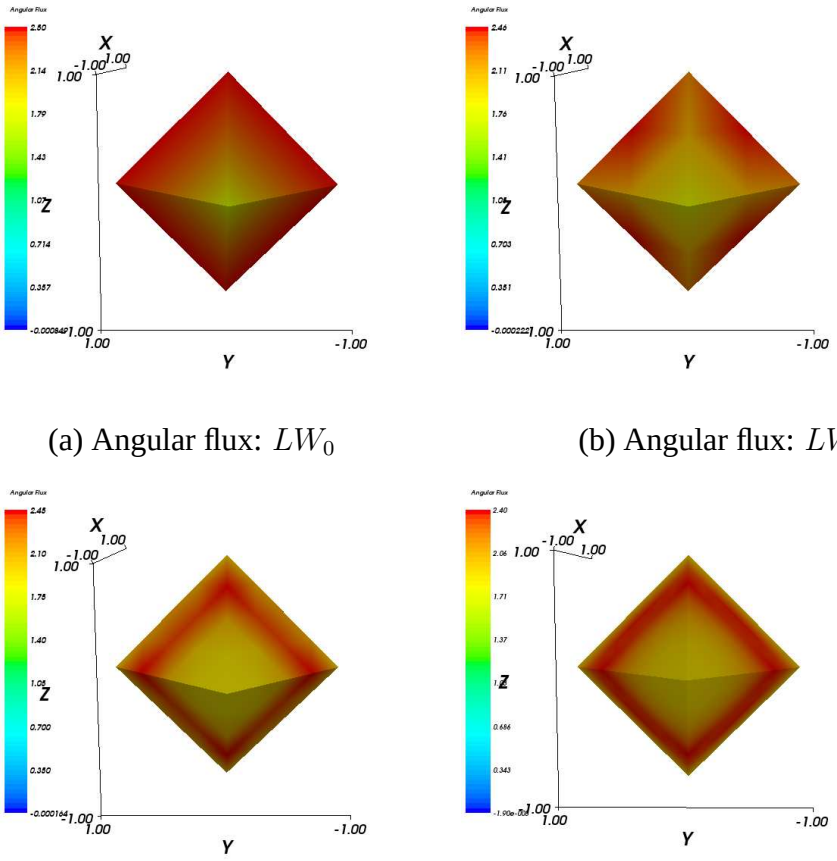

(a) Angular flux: LW_0

(b) Angular flux: LW_1

(c) Angular flux: LW_2

(d) Angular flux: LW_3

Figure 3.14: Angular flux plots at the mid point of region 3 of the Reed cell problem. The calculations were performed using the linear spherical wavelet discretisation.

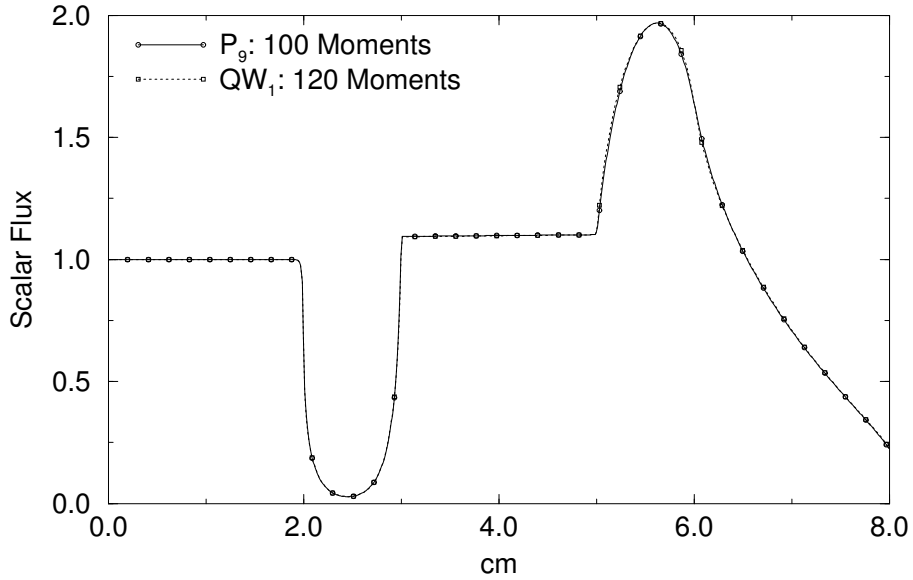

Figure 3.15: The P_9 and quadratic spherical wavelet scalar flux solutions of the Reed cell problem.

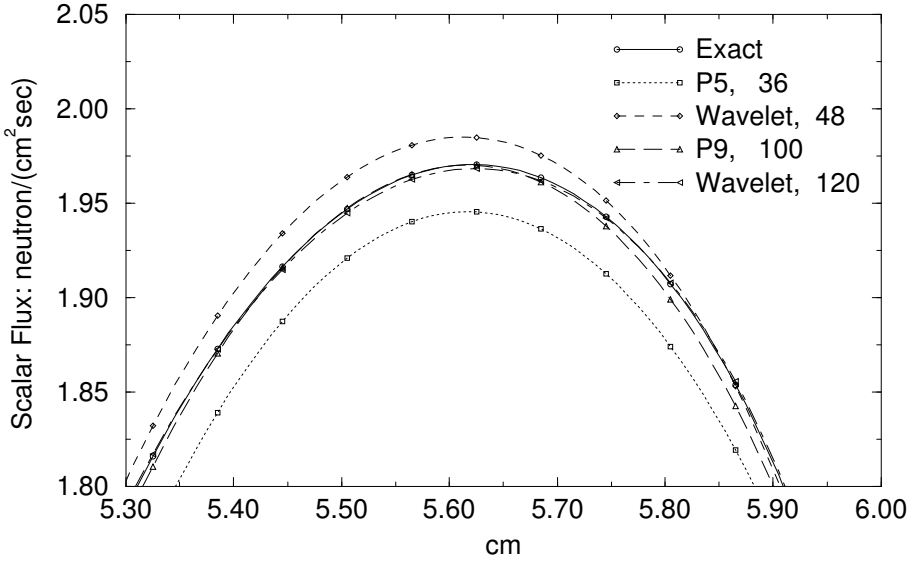

Figure 3.16: The QW_0, QW_1, P_5 and P_9 scalar flux solutions of the Reed cell problem. The solutions are concentrated over the moderator in region 4.

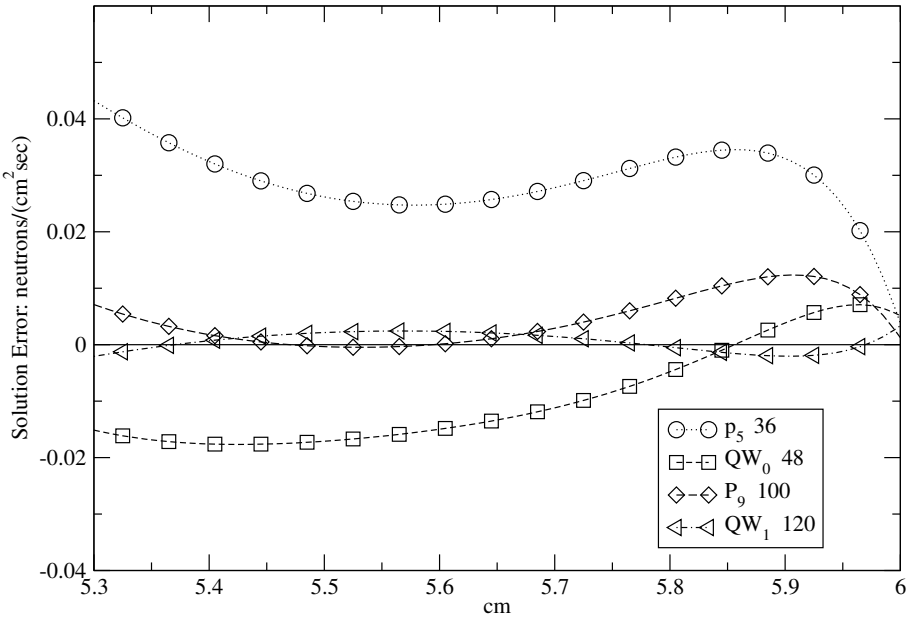

Figure 3.17: The QW_0, QW_1, P_5 and P_9 scalar flux errors of the Reed cell problem. The solutions are concentrated over the moderator in region 4.

Figure 3.18 presents the LW_1, LW_2 QW_0 and QW_1 wavelet solutions over region 4 of the problem - the corresponding difference plots are presented in figure 3.19. The graphs demonstrate that the quadratic wavelets can produce more accurate scalar flux solutions when using the same number of angular basis functions. This is particularly apparent for the lower order wavelet solutions where the errors are reduced by some considerable margin. The higher order solutions also demonstrate the quadratic wavelets to be the most accurate, although both approximations are very close to the exact solution.

The final demonstration of this section investigates the fully discontinuous linear wavelets - where the approximation can be discontinuous over all element edges of the subdivided octahedron O_j. Figure 3.20 presents the LW_1 and LW_2 scalar flux solutions for the approximations that have continuity (the default wavelets) and discontinuity over the elements of O_j. The discontinuous wavelets show improvements in accuracy over the peak flux in region 4. However, the cost of this improvement is that the set of angular basis functions increased from 48 to 96 for LW_1 and 120 to 386 for LW_2. This shows the discontinuous wavelets to be highly expensive for the small gain in accuracy. However,

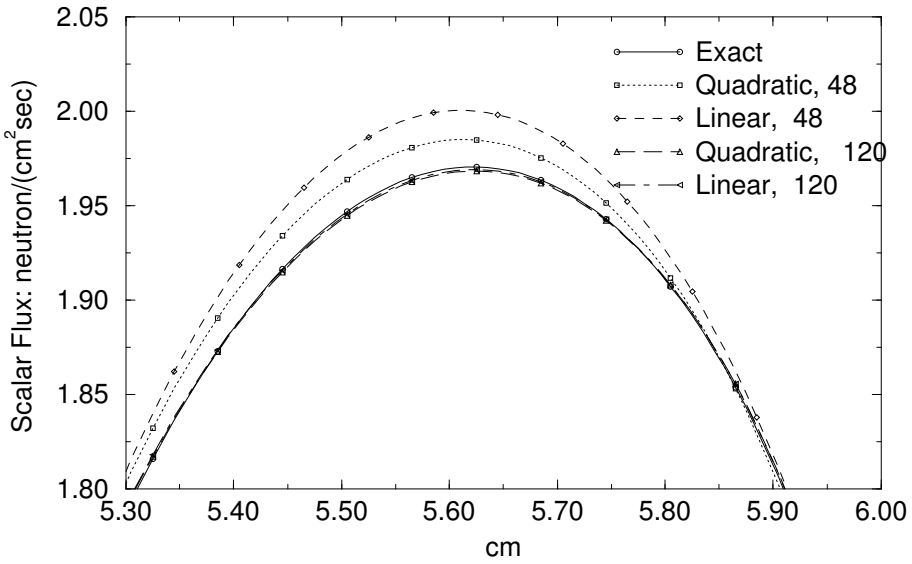

Figure 3.18: The QW_0, QW_1, LW_1 and LW_2 scalar flux solutions of the Reed cell problem. The solutions are concentrated over the moderator region 4.

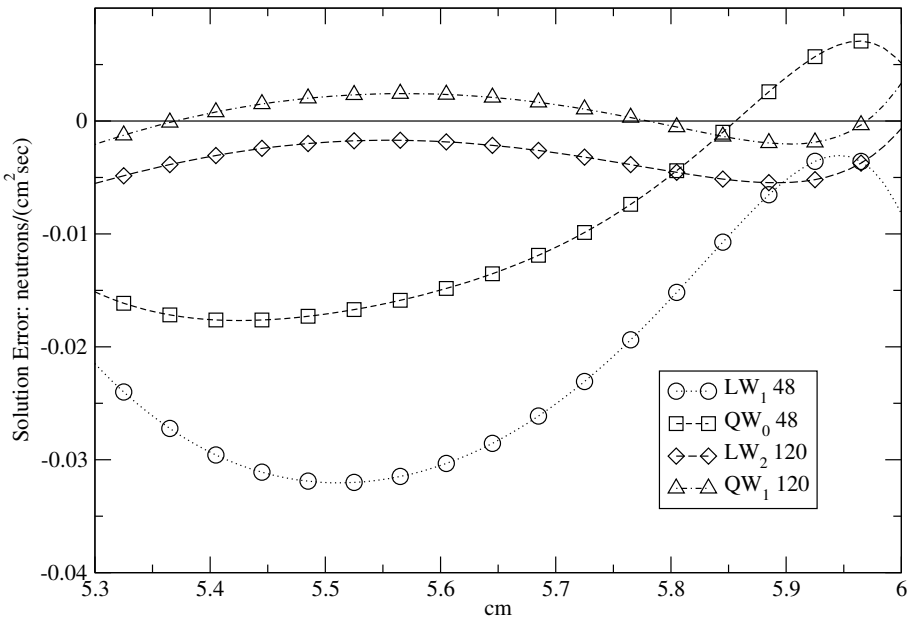

Figure 3.19: The QW_0, QW_1, LW_1 and LW_2 scalar flux errors of the Reed cell problem. The solutions are concentrated over the moderator region 4.

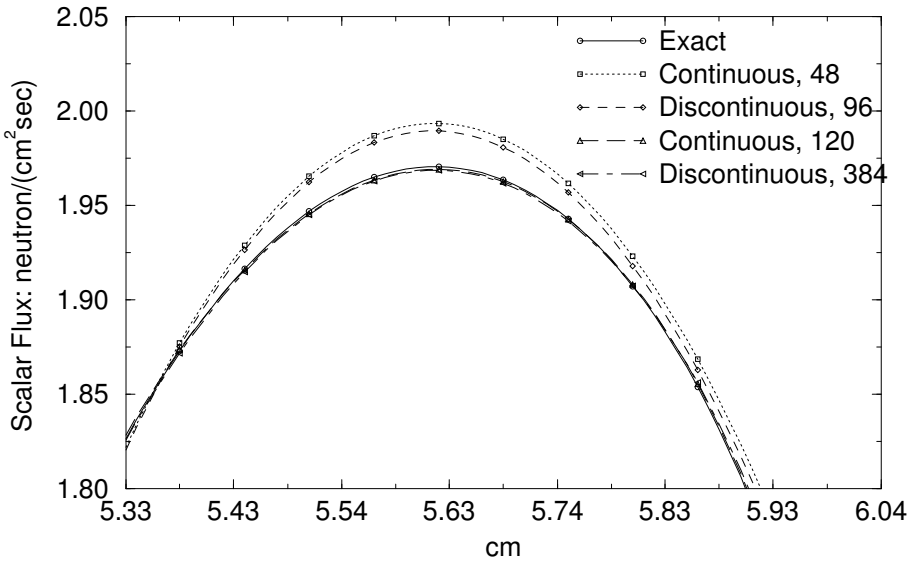

Figure 3.20: The scalar flux solution of the linear wavelets with and without discontinuity across the subdivided faces of the octahedron. The solutions are concentrated over the moderator of region 4 in the Reed cell problem.

they may yet be useful for problems where the flux also develops high discontinuities in it angular domain.

3.4.2 2-D Numerical Examples

The following 2-D numerical examples have been performed using the recently developed radiation code RADIANT. The spatial variables are discretised using the CVFEM method [88] and the Riemann approach described in section 2.3.3 is employed to resolve bare surfaces. In order to give consistent comparisons against established angular discretisation methods, S_N and P_N solutions have also been calculated using RADIANT. These solutions are used in conjunction with benchmarked solutions provided by established numerical codes.

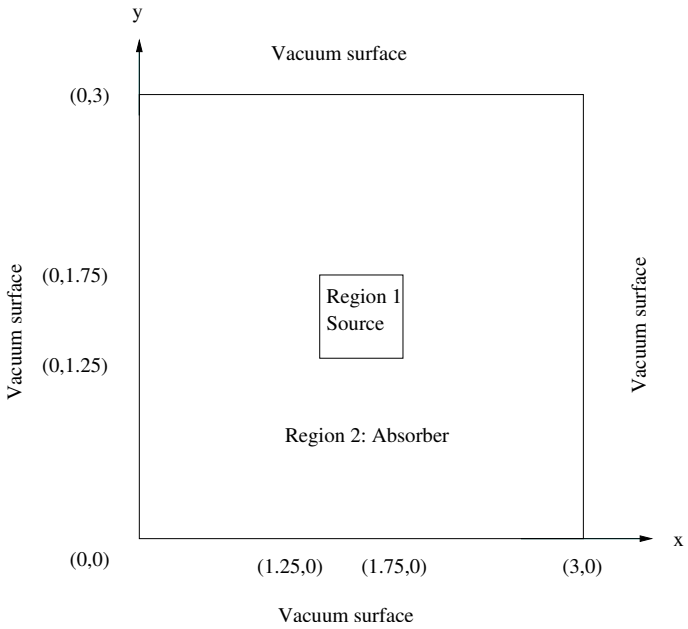

Figure 3.21: Problem 3: Diagram of the 2-D Box Source Problem.

Problem 3: A Pure Absorbing Box Problem

This numerical example is designed to demonstrate the accuracy of wavelets when solving two dimensional problem domains. The aims are compare the accuracy of the linear and quadratic wavelets and to investigate their capabilities in comparison to the conventional angular expansions.

A diagram of the problem domain is presented in figure 3.21 and the corresponding material and source information is listed in table 3.3. The problem domain is a 3cm $\times$ 3cm square which contains a material with an absorbing cross-section $\sigma_a = 0.5\mathrm{cm}^{-1}$. An isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ occupies the central square region of size 0.5cm $\times$ 0.5cm. Vacuum boundaries are prescribed along all edges and the spatial domain is discretised using a mesh consisting of 225 elements and 256 nodes.

Figure 3.22 presents the RADIANT LW_1, S_6 and P_5 scalar flux solutions. In two dimensional geometry, these approximations use 21, 24, and 24 angular basis functions, respectively. The benchmark solution provided by EVENT P_{23} is presented in figure 3.23.

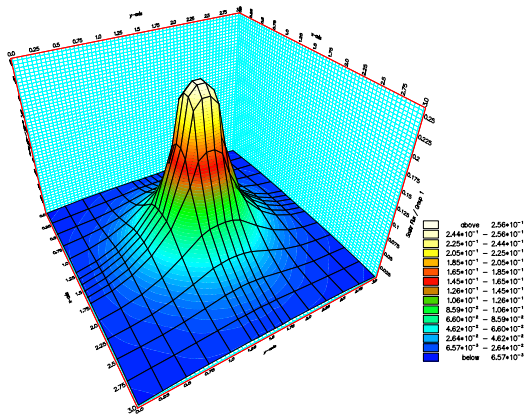

(a) P_5 solution

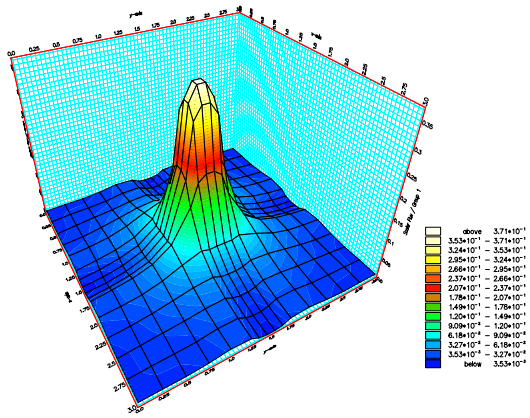

(b) S_6 solution

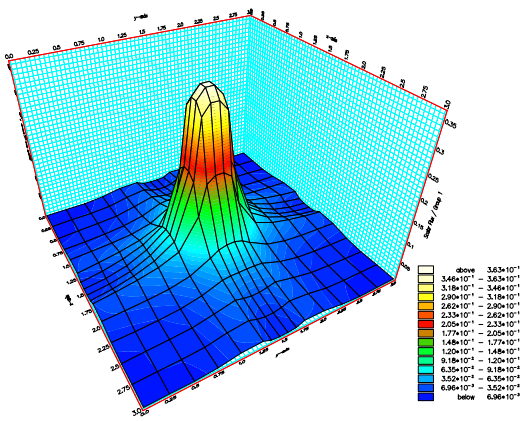

(c) LW_1 solution

Figure 3.22: The P_5 (a), S_6 (b) and LW_1 (c) scalar flux solutions of problem 3.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.0
2	0.0	0.5	0.0

Table 3.3: Source and material properties of Problem 3: The Box Source Problem 1.

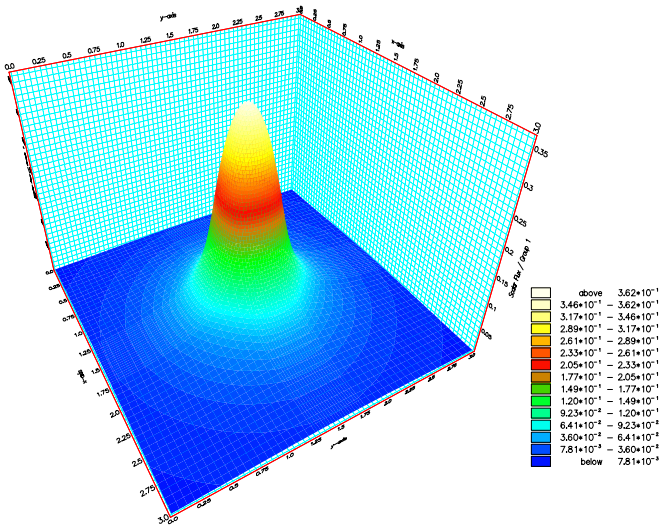

Figure 3.23: The EVENT P_{23} scalar flux (benchmarked) solution of problem 3.

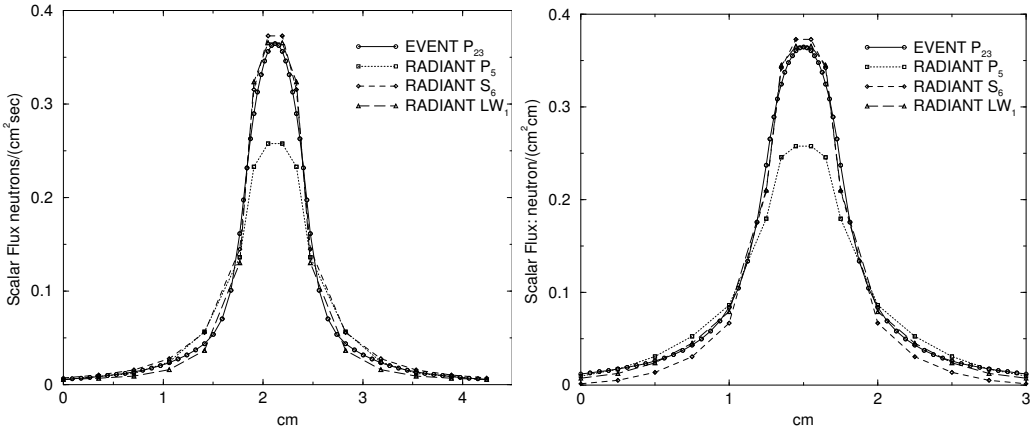

Figure 3.24: The EVENT P_{23}, RADIANT P_5, RADIANT S_6 and RADIANT LW_1 scalar flux solutions along the line (0,0) to (3,3) (left) and (0,1.5) to (3,1.5) (right) of problem 3.

Initial observation of the scalar flux profiles immediately highlights a number of effects that variation of the angular discretisation method has on the approximated solution. First it is seen that the wavelet and discrete ordinate solutions are more accurately resolved around the central region of the problem where the peak in the scalar flux occurs. The EVENT P_{23} benchmarked solution shows the scalar flux should peak at 0.34 particles.$\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This is in close agreement with the wavelet and discrete ordinate solutions, where the peak in the scalar flux reaches 0.36 and 0.37 particles.$\mathrm{cm}^{-2}\mathrm{s}^{-1}$, respectively. The P_5 solution peaks some way short of this figure at approximately 2.6 particles.$\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This is illustrated more clearly by the linear cut plots presented in figure 3.24. The two graphs, measured from points (0,0) to (3,3) and (0,1.5) to (3,1.5), confirm the P_5 approximation has indeed fallen far short of the true solution in comparison to the discrete ordinate and wavelet approximations. The peak percentage error of the P_5, LW_1 and S_6 approximation are 25.7, 2.9 and 5.7 percent respectively.

Although the last observation favours the discrete ordinate and wavelet approximations, their solutions do however contain unphysical oscillations (ray effects) in the scalar flux. Ray effects are well documented for distorting S_N solutions but they should also be expected to form in the wavelet solutions due to the absence of rotational angular invariance - the condition required for a ray effect free solution. The remainder of this numerical example will now focus on the damaging effects ray effects have on the wavelet solutions.

This demonstration compares the LW_0, LW_1 and LW_2 solutions with those of S_4, S_6 and S_{10}. Figure 3.25 presents the contours of the scalar flux solutions using all six expansion as well as the benchmarked solution and the solution using S_2.

The benchmarked solution shows the scalar flux peaks in the centre of the domain and declines uniformly from the centre towards the problem's boundaries as particles attenuate within the absorber. The S_2 approximation exemplifies how ray effects can distort the discrete ordinate solutions. The method has allowed particles to travel in the 4 directions defined by S_2 and due to the lack of scattering, the positions out of direct line with the source, with respect with the S_N directions, do not experience flux.

The S_4 and LW_0 solutions also display severe oscillations in the scalar flux. It appears that the wavelet resolve the central region outside the source to some degree better than S_4. However, both solutions are subject to large oscillation at the boundaries of the problem. The severity of the solution's oscillations are dampened by increasing the approximations to S_6 and LW_1. Visible ray effects still persist but there are indications that the wavelets are beginning to resolve the flux more accurately than S_N. Around the central part of the problem, the wavelet solution's contours resemble closely the contours of the benchmarked solution - this is something the S_6 has not been able to achieve. Increasing the approximations again to S_{10} and LW_2 removes virtually all visible ray effects. Again, it appears that the wavelet solution has a closer contour shape to the benchmarked solution than S_N. In fact, the scalar flux contours from LW_2 and EVENT P_{23} are in very close agreement.

Figure 3.26 presents the scalar flux contours of the LW_1, LW_2, QW_0 and QW_1 solutions. Although the linear and quadratic wavelet solutions are shown to vary, they do appear to be of similar accuracy. Again, there are signs of ray effects distorting the solutions. The oscillations are particularly apparent in the lower order approximations.

Problem 4: Box Problem with and without Isotropic Scattering

This numerical example is designed to demonstrate the wavelets capabilities (including investigating their susceptibility to forming ray effects) in resolving problems containing

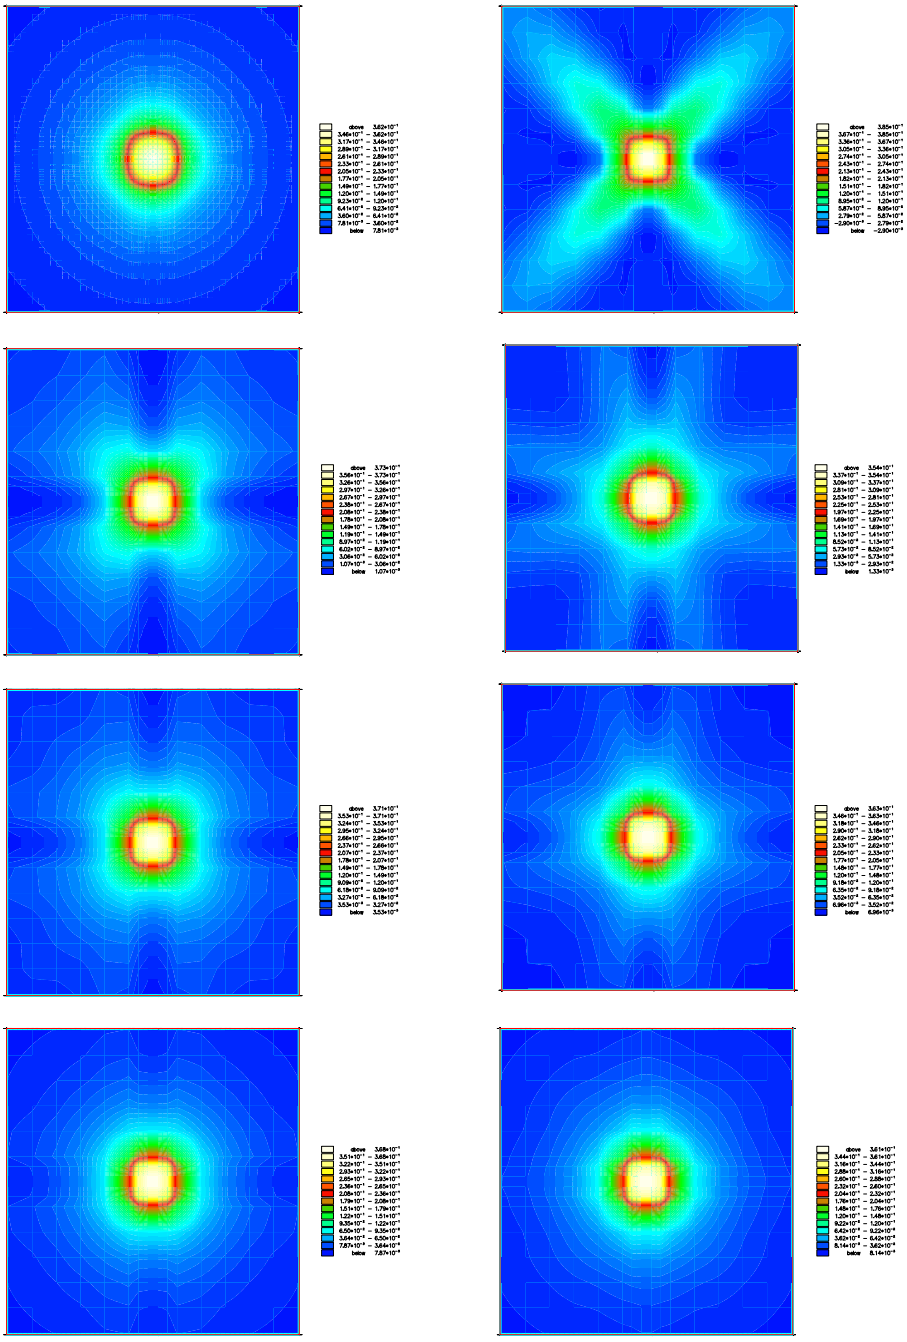

Figure 3.25: Scalar flux contour plots using S_N, P_N and LW_N angular approximations. All calculations were performed with RADIANT unless otherwise stated. From left to right, top to bottom: EVENT P_{23}, S_2 (4), S_4 (12), LW_0 (12), S_6 (24), LW_1 (24), S_{10} (60), LW_2 (60). The (.) denotes the number of angular basis functions used in the angular expansion.

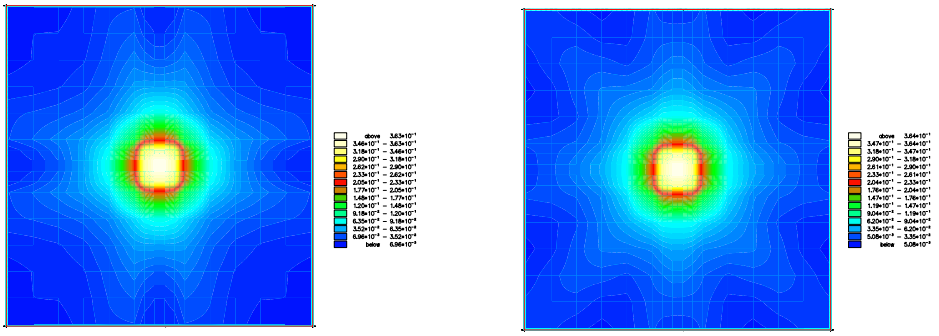

(a) LW_1 solution

(b) QW_0 solution

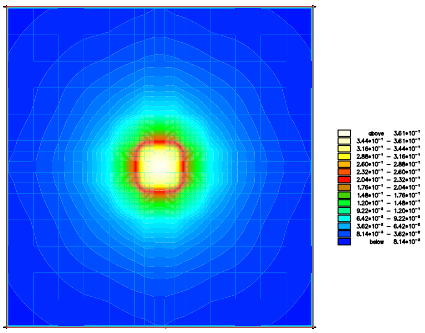

(c) LW_2 solution

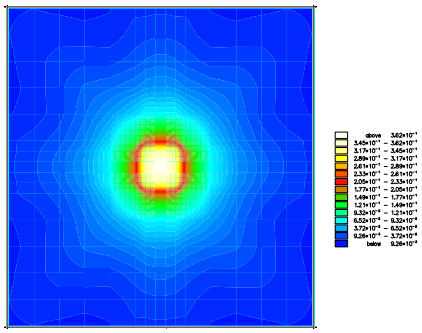

(d) QW_1 solution

Figure 3.26: This figure shows the scalar flux contours comparing the linear and quadratic wavelet solutions for problem 3: (a) LW_1 (24), (b) QW_0 (24), (c) LW_2 (60), (d) QW_1 (60). The (.) denotes the number of angular basis functions used in the angular expansion. All calculations were performed with RADIANT.

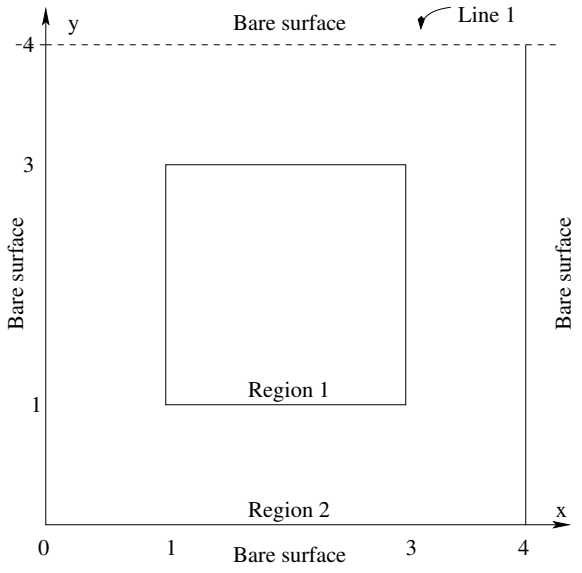

Figure 3.27: Problem 4: Diagram of the 2-D Box Source Problem.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	1.0 / 0.6	0.0 / 0.4
2	0.0	1.0 / 0.6	0.0 / 0.4

Table 3.4: Source and material properties of Problem 4.

scattering cross-sections. The default linear and quadratic wavelets are compared against the conventional P_N and S_N solutions.

The problem domain is presented in figure 3.27 for which the corresponding material and source information is listed in table 3.4. The problem under consideration frequently arises in radiation transport calculations, particularly when distortions due to ray effects are of high interest. The problem was applied in Lathrop's analysis of 2-D S_N [72] and used again in Seed's work on Walsh function approximations [77].

The problem domain is a 4cm $\times$ 4cm square containing a material with a total cross-section $\sigma_t = 1.0\mathrm{cm}^{-1}$ and a spatially constant isotropic scattering cross-section (which is varied throughout testing). An isotropic unit source of 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ occupies the central 2.0cm $\times$ 2.0cm square region and vacuum boundaries are prescribed along

all edges of the problem domain. The following calculations were performed using a bi-linear CVFEM spatial approximation over a regular grid comprising of 400 elements. A Benchmark solution was provided by EVENT P_{23}.

For this demonstration the flux leakages through the top edge boundary are examined - see line 1 of figure 3.27. Figure 3.28 presents the graphs of the LW_0, LW_1, LW_2, S_4, S_6 and S_{10} scalar flux solutions when the problem's scattering cross-section was set to zero. The results show S_4 to be more accurate in resolving the central peak region than LW_0. However, this discrete ordinate solution has also developed much larger ray-effect oscillations, which continues on previous numerical example's findings in that S_N is more susceptible to ray effects than the wavelets. The graph showing the LW_1 and S_6 solutions continue along this pattern. The results reveal S_6 to be more accurate over the central region but to still contain visible oscillations caused by ray effects. In contrast to this, the ray effects of LW_0 have been almost mitigated by LW_1. This again strengthens the view that the wavelets solutions are less susceptible to ray effects. The final graph shows the LW_2 and S_{10} solutions have both converged on the benchmark solution and contain no signs of ray effects. Figure 3.29 presents the percentage error in neutron leakage in the wavelet solutions. Over the range measured, the error of LW_1 remains below 9%, for LW_2 the maximum error reduces to 2%.

The S_N and wavelet calculations were also performed with the problem's scattering cross section set to $\sigma_s = 0.4cm^{-1}$. Figure 3.30 presents the scalar flux leakages through the top boundary. These results share close similarities with the $\sigma_s = 0$ calculations, however, ray-effects are less significant in the presence of scattering. In addition to this the LW_1 and S_6 solutions are much more accurate when scattering cross-sections are increased - this comes as a direct consequence of isotropic fluxes being induced by the scattering cross-section. Figure 3.29 presents the percentage error of the flux leakage across the top boundary. The LW_1 and LW_2 errors do not exceed 9% and 2%, respectively.

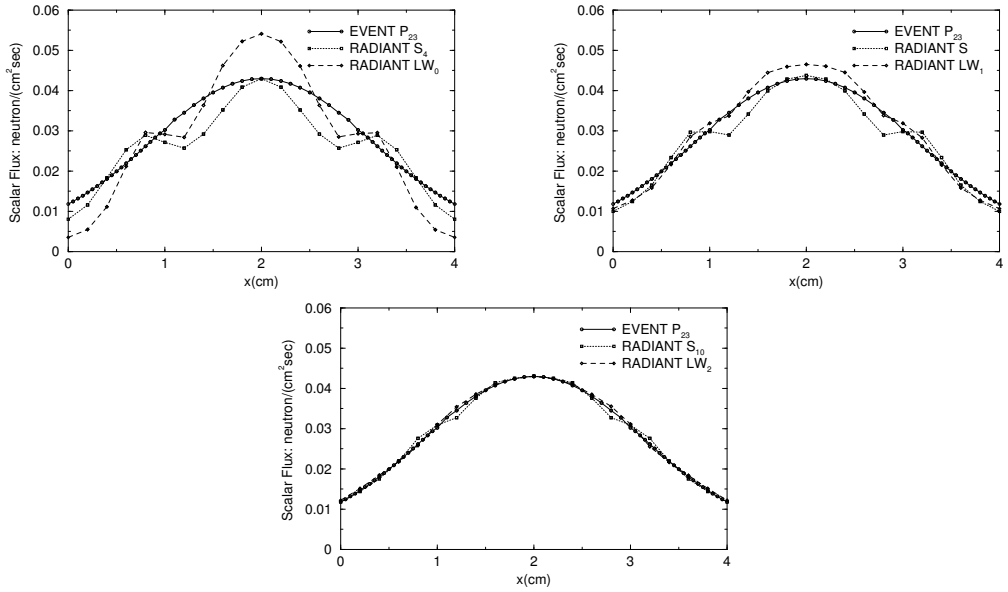

Figure 3.28: Problem 4, $\sigma_s = 0$: Scalar flux leakage through the boundary (0,4) - (4,4). Top left: S_4 and LW_0 discretisations, Top right: S_6 and LW_1 discretisations, Bottom centre: S_{10} and LW_2 discretisations.

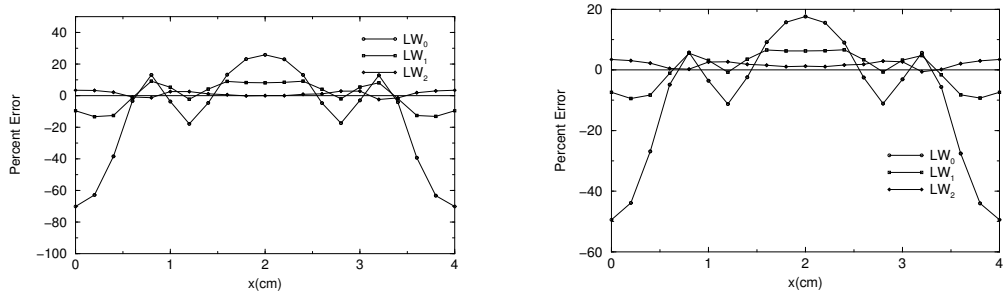

Figure 3.29: Problem 4, The percentage error of the flux leakage through top boundary using linear wavelet discretisations. Left: $\sigma_s = 0$, Right $\sigma_s = 0.4$.

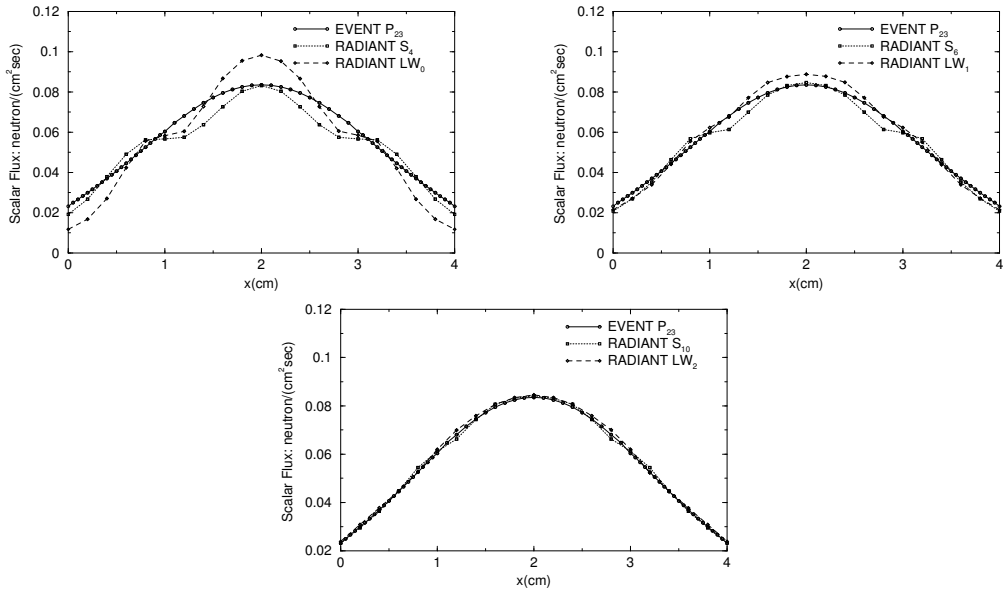

Figure 3.30: Problem 4, $\sigma_s = 0.4$: Scalar flux leakage through the boundary (0,4) - (4,4). Top left: S_4 and LW_0 discretisations, Top right: S_6 and LW_1 discretisations, Bottom centre: S_{10} and LW_2 discretisations.

Region	Source $(cm^{-2}s^{-1})$	$\sigma_a\ (cm^{-1})$	$\sigma_s\ (cm^{-1})$
1	1.0	0.5	0.0
2	0.0	0.5	0.0
3	0.0	0.0	0.0

Table 3.5: Source and material properties of Problem 5: The Straight Duct Problem.

Problem 5: Straight Duct Problem surrounded by a Purely Absorbing Region

This problem is designed to demonstrate the capabilities of the wavelet expansions to accurately resolve problems with more demanding, non uniform domains containing void material cross-sections. In this example only the default linear wavelets are considered and solutions are compared to those of the established methods (EVENT and RADIANT P_N).

The problem is the straight duct problem referenced in [34] page 449. The problem

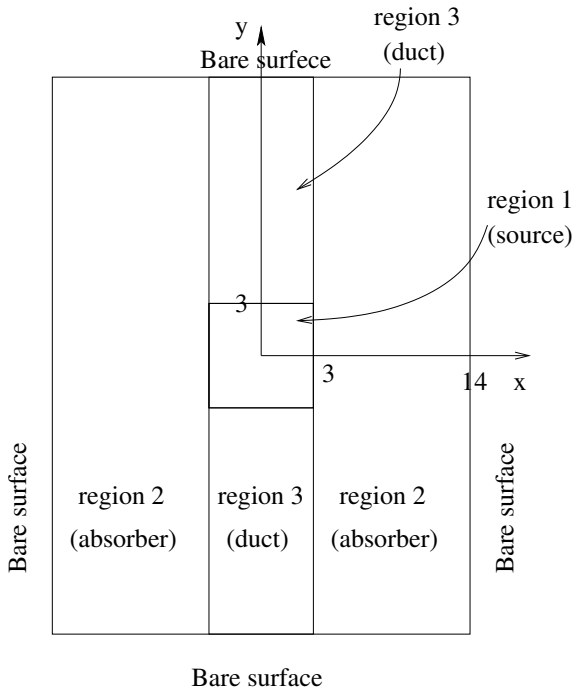

Figure 3.31: Problem 5: Diagram of the 2-D straight duct problem domain.

domain is presented in figure 3.31 and the corresponding source and material information is listed in table 3.5. The numerical example has a rectangular domain of size 36cm $\times$ 28cm. The central 6cm $\times$ 6cm square region contains an isotropic source of strength 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and a purely absorbing material with a cross-section $\sigma_t = 0.5\mathrm{cm}^{-1}$. Two ducts containing voids of width of 6cm extend from the source to the problem's upper and lower boundaries. The material surrounding the ducts has a absorbing cross-section $\sigma_t = 0.5\mathrm{cm}^{-1}$. Vacuum boundaries are prescribed along all edges of the problem. A finite element mesh consisting of 3,600 elements was employed to discretise the problem's spatial domain.

Figure 3.32 presents the EVENT P_{13}, RADIANT LW_1 and RADIANT LW_2 scalar flux solutions, all three results show close similarities in their profiles. The scalar flux solution peaks in the problem's centre at approximately 1.7 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. In the absorber region adjacent to the source the high attenuation of particles have caused the sharp decline in the scalar flux profiles. Through the ducts, for which there is no particle attenuation, the scalar flux descent is at a much lower gradient.

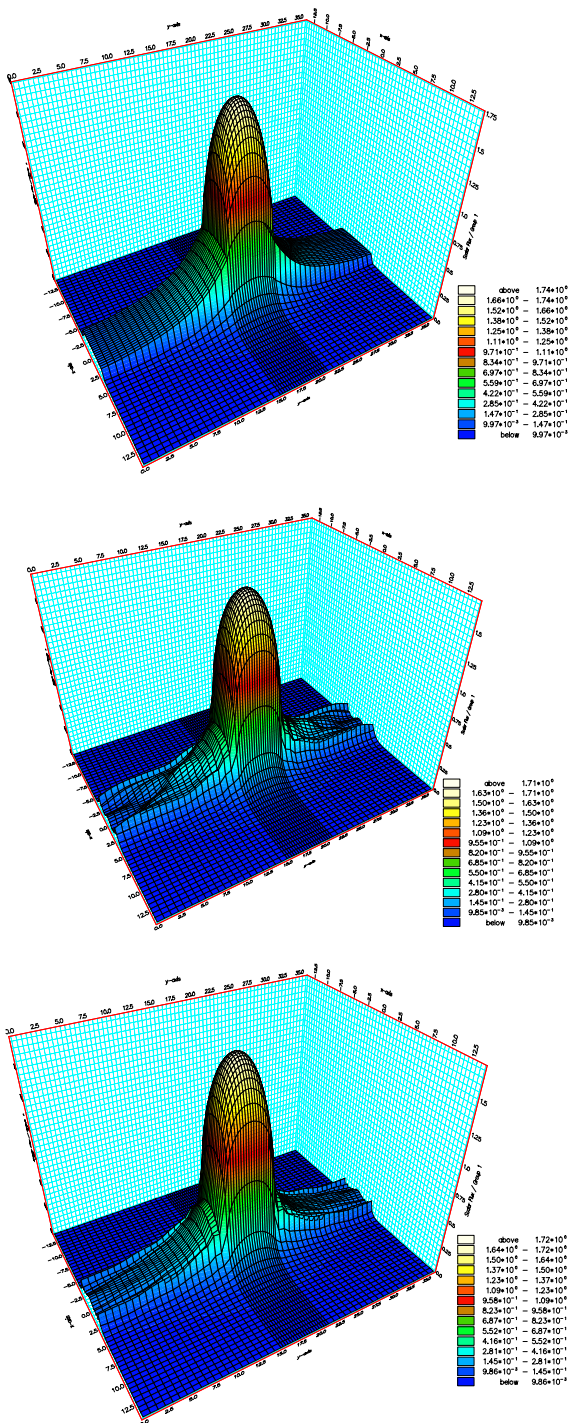

Figure 3.32: Problem 5: Scalar flux solutions of the EVENT P_{13} (top), RADIANT LW_1 (middle) and RADIANT LW_2 (bottom) discretisations.

The linear wavelet scalar flux profiles display a number of oscillations through the duct regions that are not present in the EVENT P_{13} solution. These oscillations are most likely to be unphysical fluctuations caused by either ray effects through the angular discretisation or Gibbs oscillations that form from RADIANT's spatial approximation. In order to investigate these oscillations, figure 3.33 presents the scalar flux plotted through the points from (0,0) to (0,18). Included in this graph are the DG-DO results using 2,500 directions on a fine spatial mesh - which is used as the benchmark solution. The graph indicates the oscillations to be a product of both the angular and spatial approximations. Across the interval (0,6) to (0,18), the large oscillations manifesting in the low order wavelet approximation are mitigated using the higher order wavelet expansion. This is highlighted by the difference plot presented in figure 3.34, which shows the extent to which the low order approximation's errors have been reduced. This indicates the angular approximation to be the cause of the oscillations - with ray effects being the most probable cause. However, the fluctuations over the interval (0,4) to (0,6) are not mitigated by increasing the angular approximation. In fact, the wavelet solutions are close to convergence across this interval, indicating that the oscillations are spatial discretisation errors. It is also interesting to note that the largest variation between the EVENT and RADIANT solutions occur at the surface between the duct and source regions. However, the graphs of figures 3.33 and 3.34 show EVENT P_N to be the most erroneous at these regions.

Problem 6: A Cylinder Source within a Cylinder

This Problem is designed to demonstrate the wavelets capabilities in solving demanding geometries using unstructured spatial meshes. In this demonstration the linear wavelets are compared with the solutions of EVENT and MCNP.

The problem domain is displayed in figure 3.35 and the corresponding source and material information is listed in table 3.6. The problem domain consists of a circle (or infinitely long cylinder) 15cm in radius. The central circle of 1cm radius contains an isotropic source of intensity 1.0 neutrons per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The material cross-sections are set to zero across the whole domain and vacuum boundary conditions are prescribed along the domains edge. The spatial domain has been discretised using an unstructured mesh, gen-

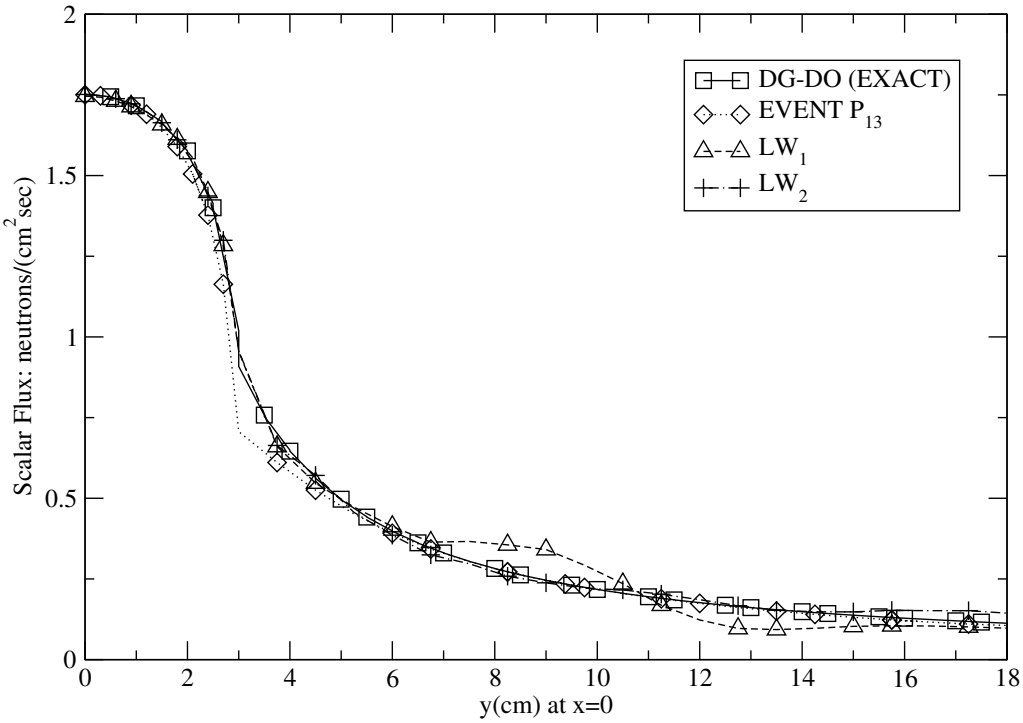

Figure 3.33: Problem 5: Scalar flux solutions of the EVENT P_{13}, RADIANT LW_1 and RADIANT LW_2 discretisations through the line (0,0),(0,18).

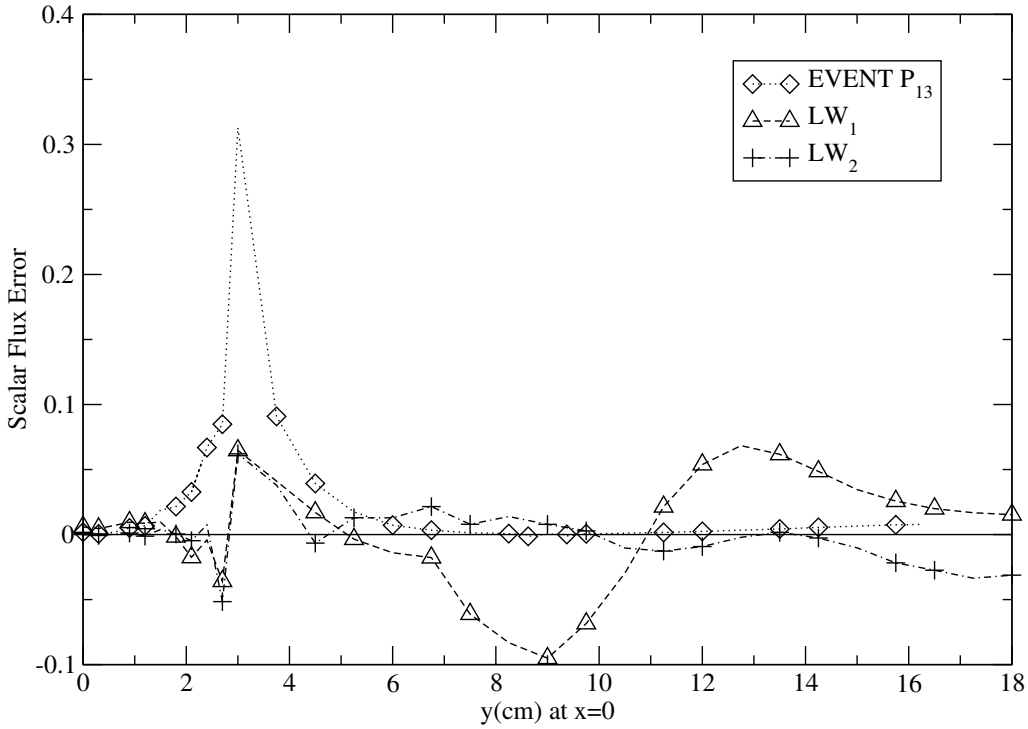

Figure 3.34: Problem 5: Scalar flux errors of the EVENT P_{13}, RADIANT LW_1 and RADIANT LW_2 discretisations through the line (0,0),(0,18).

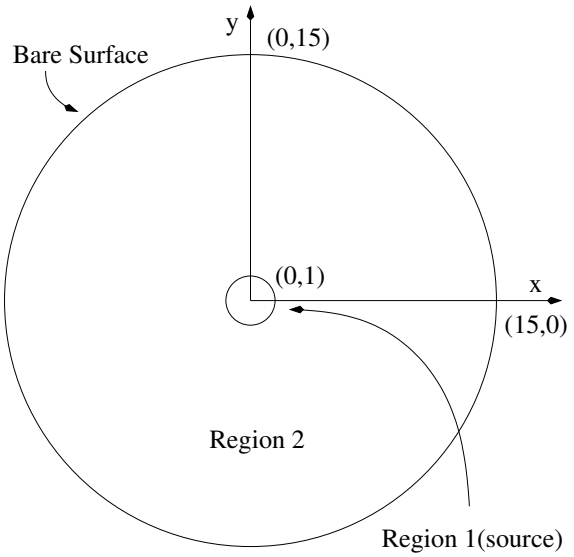

Figure 3.35: Diagram of Problem 6: A cylinder source inside a cylinder.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.0	0.0
2	0.0	0.0	0.0

Table 3.6: Source and material properties of Problem 6: The cylinder source inside a cylinder.

erated using a Delaunay triangulation [131], which consists of 8,278 triangular elements using 4,200 nodes, this is presented in figure 3.36.

For this demonstration a number of calculations were performed using P_N, S_N and linear wavelets. The EVENT P_3 and P_9 scalar flux solutions are presented in figure 3.37. For these calculations a small absorbing cross-section (0.01/cm) had been used in order to avoid break down of the even-parity method when used with voids. Figure 3.38 presents the S_2 and S_4 calculations, figure 3.39 presents the solutions of S_6 and S_{10} and figure 3.40 presents the LW_0, LW_1 and LW_2 solutions.

The solutions reveal a striking contrast between P_N, S_N and the wavelets. The P_3 solution contains a diffuse scalar flux which peaks in the central region at just under 0.2 particles

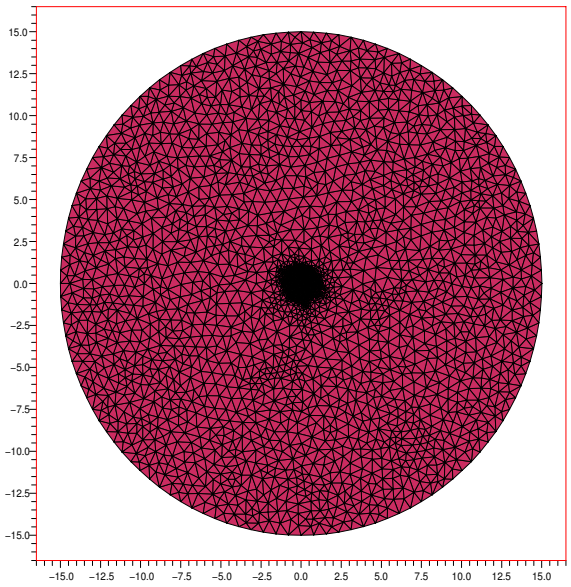

Figure 3.36: The unstructured mesh produced from a Delaunay triangulation used discretised the spatial domain of problem 6. The mesh is generated from 8,278 elements using 4,200 nodes.

per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The P_9 solution still retains a fairly diffusive scalar flux, however, the solution's profile does show a larger and more defined peak that rises above 0.35 particles per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The S_N calculations clearly conflicts with these P_N solutions, the first noticeable discrepancy being the larger peak in scalar flux over the source region. The scalar flux peaks at just over 1.4 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ for all S_N calculations, a magnitude greater than the value observed with P_N. The S_N calculations then show a very steep decline in the scalar flux profile in the outer regions of the source. On leaving the source the scalar flux either drops to zero or propagates along the directions of the S_N quadrature set. Note that once outside the source these fingers of flux remain constant in magnitude as there is no attenuation of the particles. The wavelet solutions display many close similarities to the results of S_N. The peak in scalar flux is Sharp, concentrated over the source region and reaches a magnitude of approximately 1.5 particles per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The ray effects in the solutions are also seen as fingers of flux propagating away from the source.

From these calculations it was necessary to determine which, if any, of the solutions is

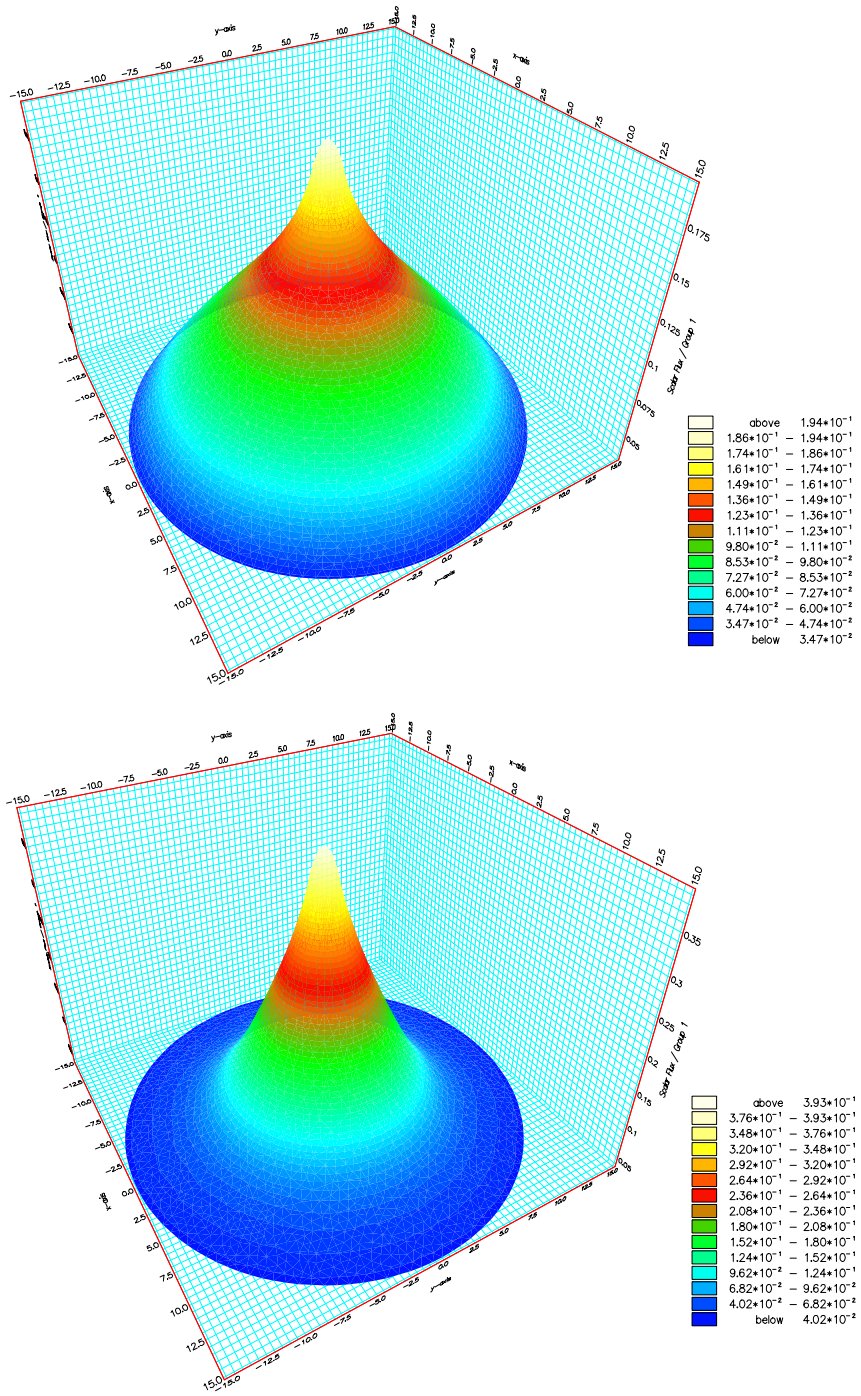

Figure 3.37: The EVENT P_3 (top) and P_9 (bottom) scalar flux solutions of Problem 6.

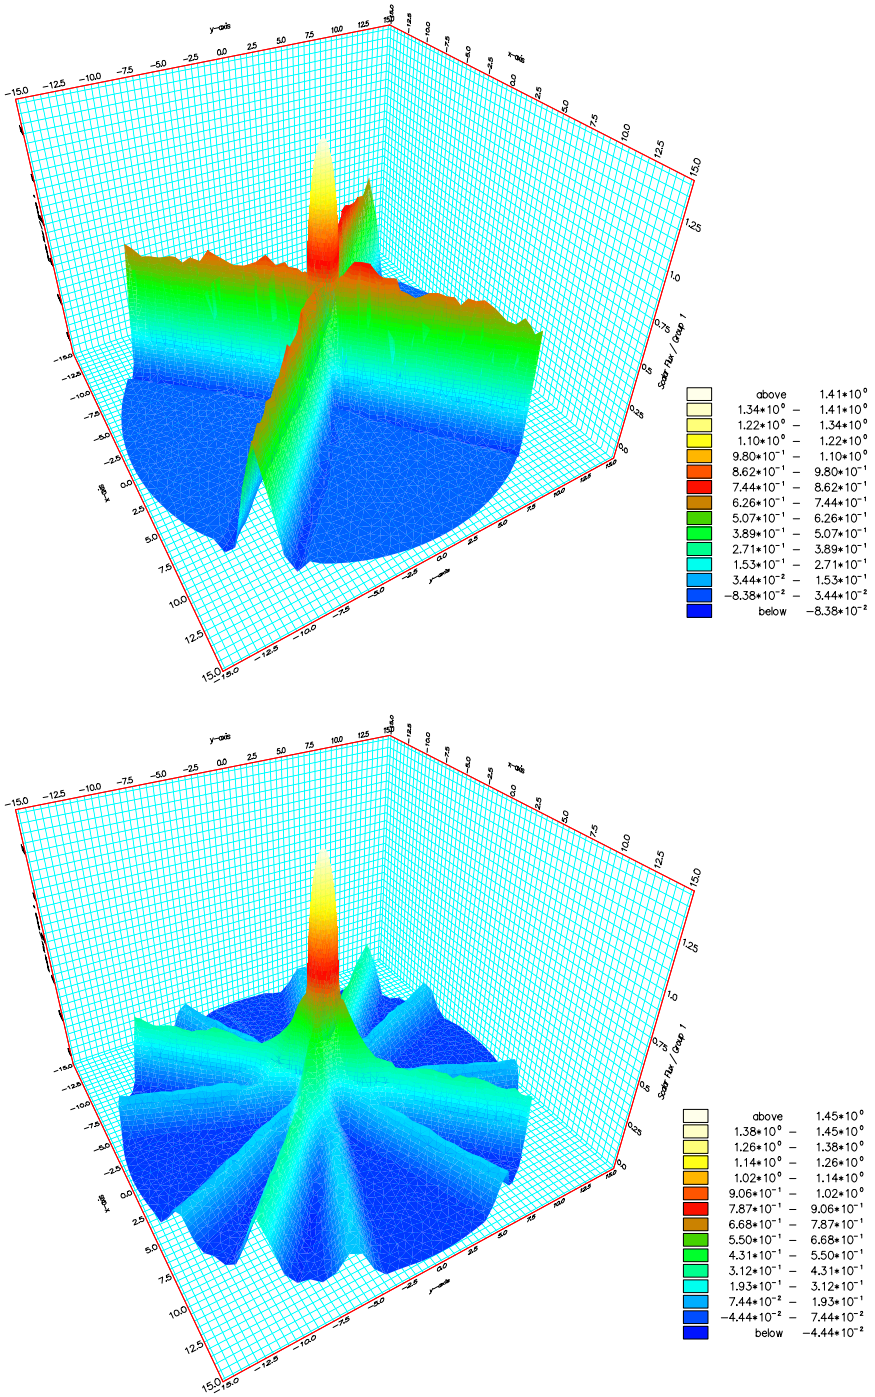

Figure 3.38: The RADIANT S_2 (top) and S_4 (bottom) scalar flux solutions of Problem 6.

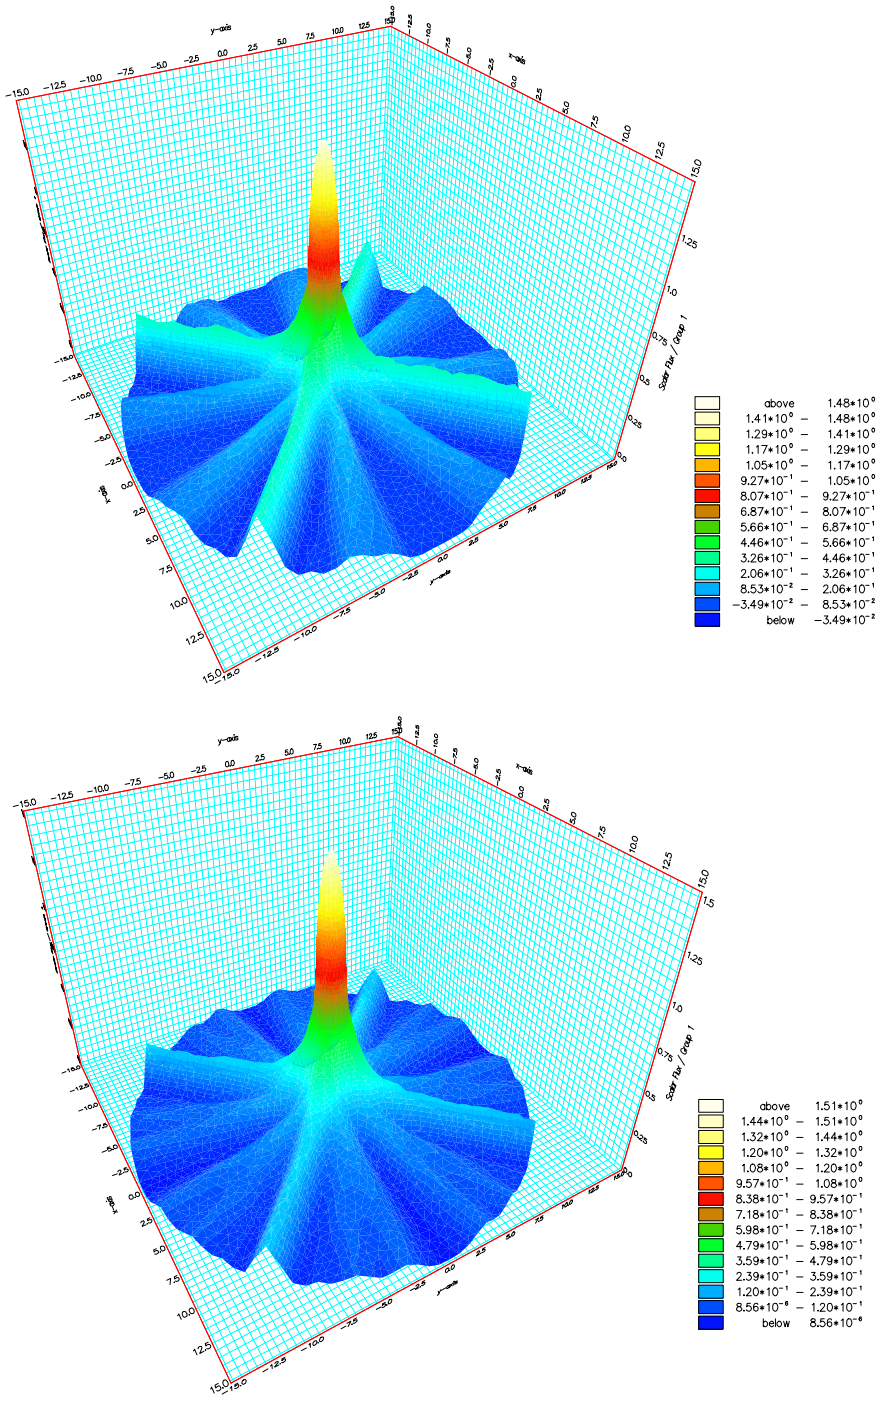

Figure 3.39: The RADIANT S_6 (top) and S_{10} (bottom) scalar flux solutions of Problem 6.

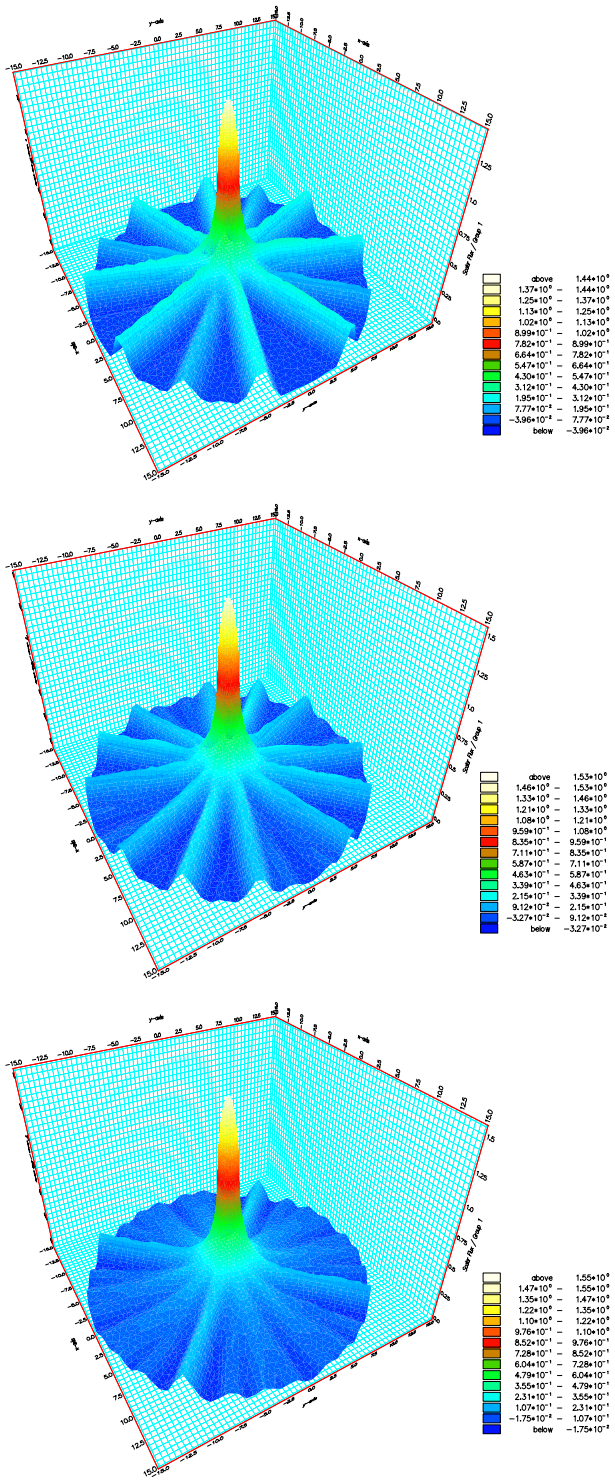

Figure 3.40: The RADIANT LW_0 (top), LW_1 (middle) and LW_2 (bottom) scalar flux solutions of Problem 6.

correct, as it is clear from these conflicting solutions that at least one of these methods is proving inaccurate within this problem domain. The first method undertaken was to introduce an absorbing cross-section into the problem's materials and investigate what effects its magnitude has on the S_N and P_N solutions (the wavelets have been omitted from this experiment due to their close agreement with S_N). Calculations were performed with S_2 (RADIANT) and P_9 (EVENT) angular approximations. The values of absorbing cross-section used were 1.0 cm^{-1}, 0.5cm^{-1}, 0.3cm^{-1}, 0.2cm^{-1}, and 0.1cm^{-1}.

The graphs presented in figure 3.41 display the flux profiles along the line through the points (-15,-15) to (15,15). The top graph of this figure contains the S_2 calculations. These calculations show what would appear to be an expected flux profile for the various absorption cross-sections. That is, a higher peak in the scalar flux is gained upon a reduction in the attenuation of particles. The bottom graph of figure 3.41 contains the P_9 calculations. For the calculations with cross-sections exceeding 0.3 cm^{-1}, the peaks in scalar flux are in close agreement with those of S_2. However, any further reductions in the cross-section resulted in a reduced peak in the scalar flux, this is clearly incorrect.

From this observation it appears that P_N is showing inaccuracies in the resolution of the domain as the material cross-section approaches zero. A possible cause for these inaccuracies may lie in this method's inability to resolve the angular flux in the outer regions of the problem. In these regions the angular flux distribution is highly peaked, due to the distance from the source and the lack of attenuation and scattering of particles, and is therefore a demanding distribution to approximate. As the P_N method finds its best fit solution, large errors occurring in these outer regions could effect the solutions in other parts of the domain. If this is to be the case then the effects could account for the observed reduction in the scalar flux.

It is possible to test this conjecture by modelling the source region with the outer disk removed. This is due to the inner disk's independence from the outer disk when no scattering is present. Calculations were performed with and without the outer disk using EVENT P_1, P_9 and P_{19}, with the material cross-section set to 0.1 cm^{-1}. Figure 3.42 presents the scalar flux along the line at the points (-1,-1) to (1,1) for the source only calculations and (-15,-15) to (15,15) for the calculations using the full domain. These graphs

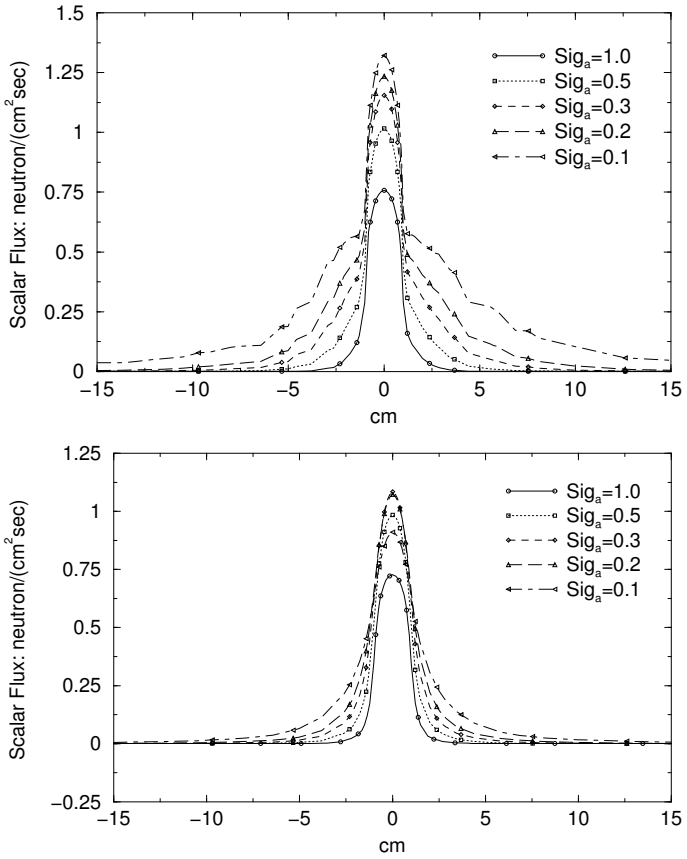

Figure 3.41: The scalar flux solutions of problem 6 through the points (-15.-15) to (15,15). In these experiments the absorbing cross-section is varied between the values 1.0 and 0.1 cm^{-1}. Top: S_2 solutions, Bottom: P_9 solutions.

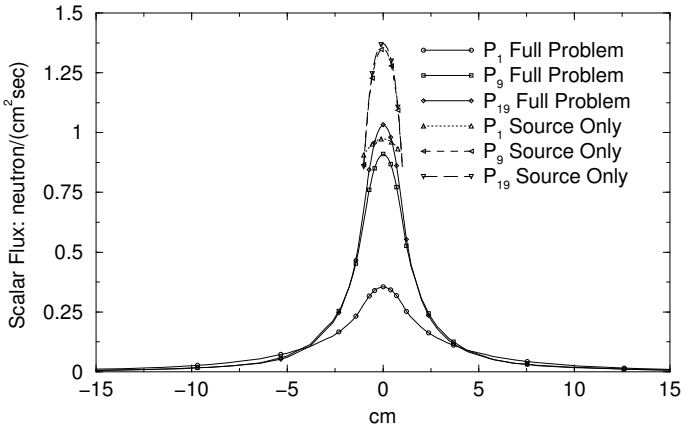

Figure 3.42: The scalar flux solutions to problem 6 with an absorbing cross-section of 0.1 cm^{-1}. Calculations were obtained using P_1, P_9 and P_{19} angular expansions. The solutions to just the source region using these angular expansions are included.

show vast increases in the scalar flux peak with the outer disk removed. The P_1 calculations show an increase from 0.35 particles per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ to 0.96 particles per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The P_9 calculations increase from 0.9 to 1.3 particles per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and the P_{19} increase from 1.0 to 1.34 particles per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This confirms the inability of P_N to model the whole domain involving low cross-section materials. Even a P_{19} approximation was unable to give an accurate approximation over the source region.

In addition to this conclusion it should also be noted that the P_9 and P_{19} calculations almost converge on the source only calculations, indicating the peak in the scalar flux should be approximately 1.3 particles per $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This is in close agreement with the S_N calculations presented in figure 3.41. This figure is also just shy of the peak flux calculated by S_N and the wavelets, figures 3.38 - 3.40. However, this small shortfall should be expected due to the introduction of the small cross-section used in the experiment. This makes it possible to conclude that the wavelet and S_N approximations of figures 3.38 - 3.40 are producing accurate representations, at least over the source region.

The accuracy of the wavelet and S_N solutions are confirmed by their comparison with MCNP calculations. Figure 3.43 presents the scalar flux profiles of the MCNP, LW_2 and P_9 solutions through the line (0,0)-(15,4). The wavelet and MCNP solutions are in almost

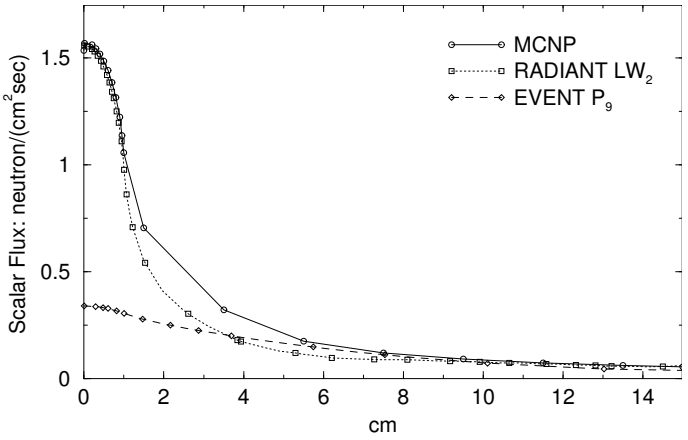

Figure 3.43: The MCNP, LW_2 and P_9 scalar flux solutions of problem 6. The line cut was taken from the problem's centre to the boundary edge.

exact agreement over the source region - with the P_9 solution far short of their maximum. Therefore the three experiments used to investigate the differences in the P_N and wavelet solutions have all agreed that the wavelet calculations are correct.

Figure 3.44 presents the contour plots of the LW_1, LW_2, S_6 and S_{10} solutions. Both contours of the S_6 and LW_1 solutions display ray effects. However, the ray effects in the LW_1 solutions only become apparent several centimetres from the source region. The area just within a few centimetres of the source region show a uniformly declining scalar flux profile from the problems' centre, as would be expected. The S_6 approximation shows ray effects dominating the solution from the source region. At no point from the centre is there a uniformly declining scalar flux.

The S_{10} solution still displays ray effects that are visible from a close distance to the central region. In addition to this there are many areas towards the outer boundary that do not experience flux. In contrast to this, the wavelet LW_2 approximation shows a uniformly declining flux several centimetres from the centre. Furthermore, although the solution are still contains visible fluctuations, far more of the outer regions experience flux. It can be concluded from these observations that the wavelets have given a far more realistic scalar flux profile than S_N. The wavelets have once again shown that their solutions are also less susceptible to ray effects than the solutions of S_N.

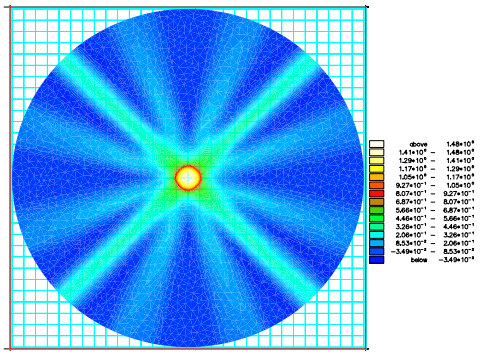

(a) S_6 (24)

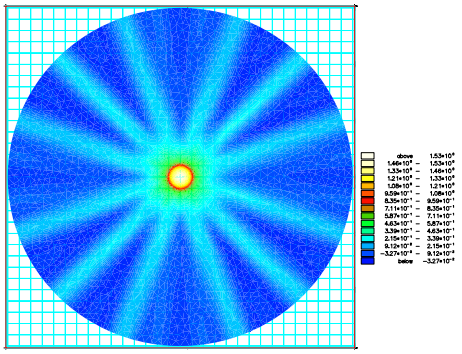

(b) LW_1 (24)

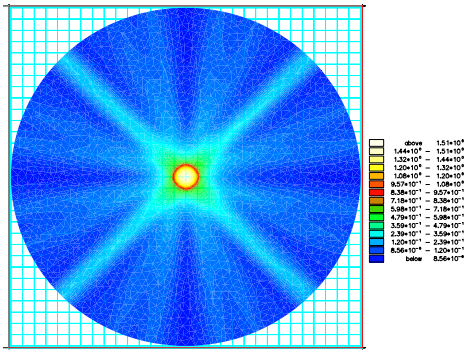

(a) S_{10} (60)

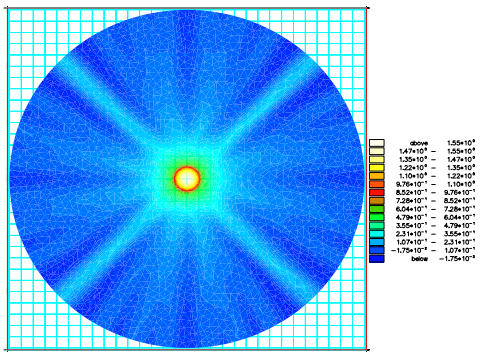

(b) LW_2 (60)

Figure 3.44: The scalar flux contours of the solutions to problem 6. (a) S_6, (b) LW_1, (c) S_{10}, (d) LW_2. The term (.) denotes the angular expansion size.

3.5 Conclusions

In this chapter two octahedral spherical wavelet bases have been developed and applied to discretise the angular variable of the Boltzmann transport equation. The wavelet functions are based on Swelden's second generation wavelets, however, in this work a modification was developed in order to mitigate the ray-effect type errors that form in the standard wavelet solutions. A number of numerical tests were performed in order to demonstrate the wavelets capabilities and to highlight their susceptibility in forming ray-effects in their solutions. Solutions were compared with the solutions of established angular discretisation methods (where the same code discretising the spatial domain was used to ensure fair comparisons) and other established radiation transport codes.

The 1-D numerical examples demonstrated the modified wavelets proposed in section 3.3.5 mitigated the ray effects type errors seen in the solutions of the original wavelet bases. In addition to this, the modified wavelets were demonstrated to be capable of resolving the demanding and complicated angular distributions generated by the Reed cell problem. It was also shown that the fully discontinuous wavelets were highly expensive, in terms of the number of angular unknowns, for the small gains in accuracy they gave.

The 2-D numerical examples highlighted a number of important issues with regards to the wavelets capabilities. First, in all examples the wavelets performed comparatively well against the established S_N and P_N discretisations. In some test cases the wavelets were shown to out performed the traditional techniques by some considerable margin. This was particularly highlighted in the final numerical example involving the void cylinder problem. In this demonstration the P_N method produced very poor quality solutions that showed very little resemblance to the exact solution. Not only did the wavelets manage to provide highly accurate solutions to this problem but they also demonstrated themselves to be less susceptible than S_N to forming ray effects in their solutions. The general pattern of the other multi-dimensional numerical examples agreed with these findings. They also demonstrated the wavelet bases to capable of resolving isotropic scattering cross-sections and highly directed fluxes that occur from neutron streaming within voids.

The first two multidimensional problems showed the linear and quadratic wavelets gave

solutions of similar accuracy - with respect to the number of angular basis functions used. However, in the one dimensional problems the quadratic wavelets showed more potential, as they clearly gave faster convergence towards the exact solution. But, in general, these results show it maybe be more beneficial to use the linear wavelets. This is due to them producing sparser angular matrices which, in turn, can be utilised to improve the efficiency of a linear solver's operations - for example, performing a matrix vector multiplication in a GMRES type algorithm.

Chapter 4

LINEAR AND QUADRATIC HEXAHEDRAL WAVELETS ON THE SPHERE FOR ANGULAR DISCRETISATIONS OF THE BOLTZMANN TRANSPORT EQUATION

Contents

In this chapter two new spherical wavelet bases for the angular discretisation of the Boltzmann transport equation are presented. The wavelet functions have been developed on a second generation wavelet framework. The bases are designed to represent spherical functions over the surface of a hexahedron by providing linear and quadratic representations over a quadrilateral partitioning of the hexahedron's surface. A demonstration of the new wavelet bases capabilities in solving neutron transport problems is presented. Four 2-D steady state transport problems are solved and the wavelet solutions are compared for accuracy and susceptibility to ray effects with the solutions of established techniques. Benchmark solutions for all transport problems are provided by established transport codes.

4.1 Introduction

In this chapter we continue to develop new numerical methods for approximating spherical functions with the intention of widening the options of angular discretisation methods for modelling particle transport problems. The intension of this chapter is to expand the class of spherical wavelet methods that was initiated in chapter 3. The wavelets previously developed were based on Swelden's spherical wavelets that approximated spherical functions over a triangular partitioning of an octahedron's surface (this surface was used to represented the surface of the unit sphere). These discretisation methods were shown to be capable of producing accurate approximations that were often more accurate than

those of the established methods S_N and P_N. In addition to this the wavelet methods were also less susceptible than S_N to forming ray effects in their scalar flux solutions.

The work in this chapter provides an extension to the methods developed in chapter 3 by developing a second low order spherical wavelet basis. In this work the spherical wavelets are designed to approximate the direction of particle travel over a quadrilateral partitioning of a hexahedron's surface, this is presented in figure 4.2. In a similar manor to the approach of chapter 3, these wavelet bases generate linear and quadratic approximations over the quadrilateral elements.

The motivation behind developing this second alternative approach is because the theory of finite element analysis generally shows quadrilateral finite elements to be more accurate and stable than triangular discretisations. These new spherical wavelets are therefore expected to inherit this property and provide a more stable and accurate approximation scheme when used to discretise the angular variable of the Boltzmann transport equation. This added stability and accuracy should also further mitigate un-physical scalar flux oscillations caused by ray effects. In addition to this, these wavelets also possess the necessary ingredients for adapting in angle. That is they are hierarchical expansions and the functions have compact support (see chapter 3), and this is utilised in a chapter 6.

In this chapter the theory of the linear and quadratic hexahedral wavelets are presented and their capabilities are demonstrated and tested against conventional methods for accuracy and susceptibility to forming ray effects. The following sections are as follows: Section 4.2 introduces the angular discretised Boltzmann transport equation. Section 4.3 defines and builds the linear and quadratic hexahedral wavelets on the sphere. In section 4.4 four 2-D steady-state neutron transport problems are solved in order to demonstrate the wavelets abilities in approximating the angular flux. This chapter is finished with a conclusion in section 4.5.

4.2 The First Order and Angularly Discretised Transport Equation

This chapter begins with the first order Boltzmann transport equation that is discretised in angle,

$$\frac{1}{v}A\frac{\partial\Psi(r,t)}{\partial t}+A_x\frac{\partial\Psi(r,t)}{\partial x}+A_y\frac{\partial\Psi(r,t)}{\partial y}+A_z\frac{\partial\Psi(r,t)}{\partial z}+\mathcal{H}\Psi(r,t)=\mathcal{S}(r,t), \quad (4.1)$$

using an arbitrary angular approximation scheme,

$$\psi(r,\hat{\Omega},t)\approx\sum_{j=1}^{\mathcal{M}}\Psi_j(r,t)\mathcal{G}_j(\hat{\Omega}). \quad (4.2)$$

The term $\Psi(r,t)$ denotes the vector of angular coefficients corresponding to expansion (4.2). The term $\mathcal{S}(r,t)$ denotes the angular discretised source and A, A_x, A_y, A_z and $\mathcal{H}$ denote the angular mass, angular Jacobians and scattering/removal matrices. All these terms are fully discussed and defined in section 2.2.

4.3 Linear and Quadratic Hexahedral Wavelets

In the following sections the linear and quadratic spherical wavelets are developed. The wavelets bases are designed to approximate spherical functions over the surface of the hexahedron H. The following sections are set out as follows. Section 4.3.1 defines a mapping for functions on the surface of the unit sphere to the hexahedron. Section 4.3.2 describes a subdivision scheme on the surface of the hexahedron. In section 4.3.3, two multiresolution analysis (MRS) are developed. These MRAs are required to develop the linear and quadratic wavelets in section 4.3.4.

4.3.1 Mapping the Angular Flux to H

In order to represent the angular flux over the hexahedron a mapping from the sphere S^2 to the hexahedron H is required. The mapping used here is defined by the normalisation

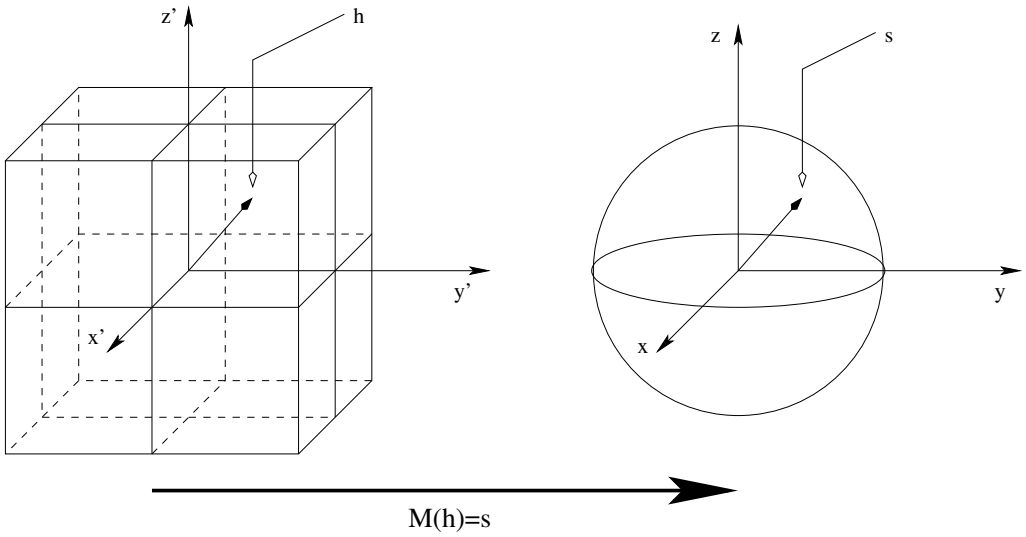

Figure 4.1: A diagram illustrating the normalising mapping from the hexahedron to the spherical domain: $M(h) = h/||h|| = s$ for $h \in \mathsf{H}, s \in S^2$. The points h and s can be expressed by vectors from the origin to the surface of their domain, that is, $h = (x^{'}, y^{'}, z^{'})$ and $s = (x, y, z)$. The relationship between the two positions are given by $s = (x, y, z) = (x^{'}, y^{'}, z^{'})/(x^{'2} + y^{'2} + z^{'2})^{\frac{1}{2}}$

of the angular vector on H, as illustrated in figure 4.1. This projection $M : \mathsf{H} \rightarrow S^2$ is defined by $M(h) = h/||h|| = s$ for $h \in \mathsf{H}$ and $s \in S^2$. This ensures the existence of a function F on H to be equivalent to any function f on S^2 through the relationship, $f(s) = F(M^{-1}(s)) = F(h)$.

In order to approximate the angular flux over H a number of modifications to the angular discretised transport equations are required. These adjustments were described in section 3.3.1 where an approximation of the octahedron was developed.

4.3.2 The Subdivision Scheme on H

Construction of the wavelets requires a subdivision scheme on the hexahedron H. The subdivision scheme used here recursively subdivides the faces of H into 4 new identical faces. The faces generated by j subdivisions are denoted by H_j. Figure 4.2 presents the

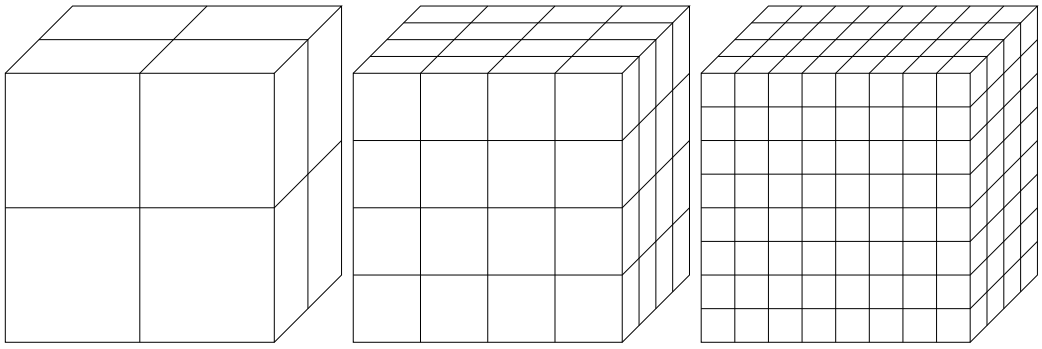

Figure 4.2: Diagram presenting the subdivision scheme employed on the faces constructing the hexahedron H. Left: The original faces, Middle: The first subdivision, Right: The second subdivision. The wavelet approximation requires discontinuity over the internal lines of the left hexahedron. Note that there can be continuity over all lines generated by the subdivision scheme and the lines defining the edges of the hexahedron.

faces H_0, H_1 and H_2.

An indexing set S_j is used to denote the vertices of the subdivided faces of H_j for the linear wavelets and the vertices, mid edge and mid face positions of the elements of H_j for the quadratic wavelets. For both wavelet methods, the definition of S_j is given as,

$$S_j = \{s_{j,k} \in \mathsf{H}_j, k \in K(j)\}, \quad j \in \{0, 1, 2, \ldots, \infty\}, \tag{4.3}$$

where $K(j)$ indexes the vertices (linear wavelets) or vertices, mid-edge and mid-element points (quadratic wavelets) of the faces H_j.

For both methods the set S_j are nested, i.e. $S_j \subset S_{j+1}$ for all $j \in \{1, 2, \ldots, \infty\}$, and so the elements in $K(j)$ are ordered to ensure the condition $s_{j,k} = s_{j+1,k}$ will always hold. The index set $M(j)$ is used to denote the positions indexed by $K(j+1)$ that are created on the $j+1$ subdivision. That is, $M(j)$ is defined as $M(j) = K(j+1) \setminus K(j)$ where $'\setminus'$ is the set removal operator. Figure 4.3 presents a diagram displaying the sets S_j, $K(j)$ and $M(j)$ over a section of the subdivided H_j.

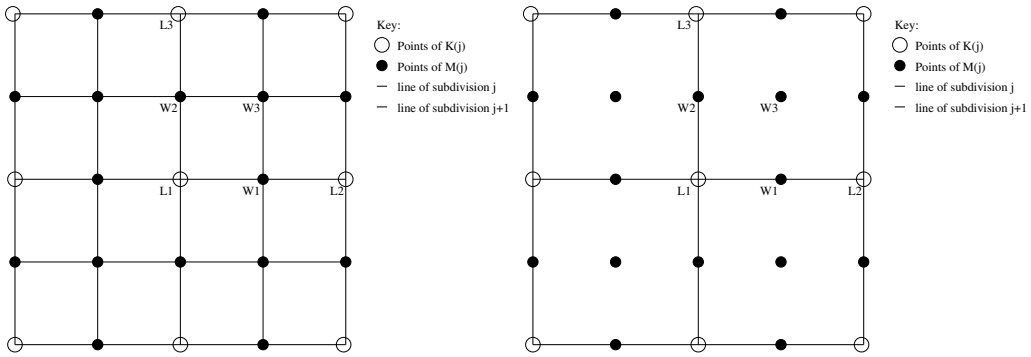

Figure 4.3: This diagram displays the points indexed by the sets $K(j)$ and $M(j)$ over a section of the hexahedron H. Left: Linear case, Right: Quadratic case.

4.3.3 A Multiresolution Analysis defined on the hexahedron H

A multiresolution analysis (MRA) defined on H is a nested sequence of closed subspaces $V_j \in L_2(\mathsf{H}), j \in \{0, 1, \ldots, \infty\}$, that satisfy the following conditions,

C1 $V_0 \subset \ldots V_{j-1} \subset V_j \subset V_{j+1} \ldots \subset L_2(\mathsf{H})$

C2 $\bigcup_{j=0}^{\infty} V_j$ is dense in $L_2(\mathsf{H})$

C3 At each level j there exists a set of scaling functions $\varphi_{j,k}, k \in K^{'}(j)$ for some index set $K^{'}$, for which the set $\{\varphi_{j,k}, \ \ k \in K^{'}(j)\}$ forms a Riesz basis of V_j.

The space $L_2(\mathsf{H})$ denotes the set of all Lebesgue functions defined on H. Two separate MRAs used in the construction of the linear and quadratic wavelets are now developed.

Linear MRA on H

The linear MRA defines the spaces V_j as the collection of all functions that are piecewise bi-linear over the elements H_j. This definition of the spaces V_j satisfy the conditions of an MRA. The first condition holds since the elements of H_j are generated from H_{j-1}, for all values of j. This therefore guarantees nesting of the functional spaces V_j. Also, in the limit of $j \to \infty$ the elements of H_j become arbitrarily small. Therefore a function

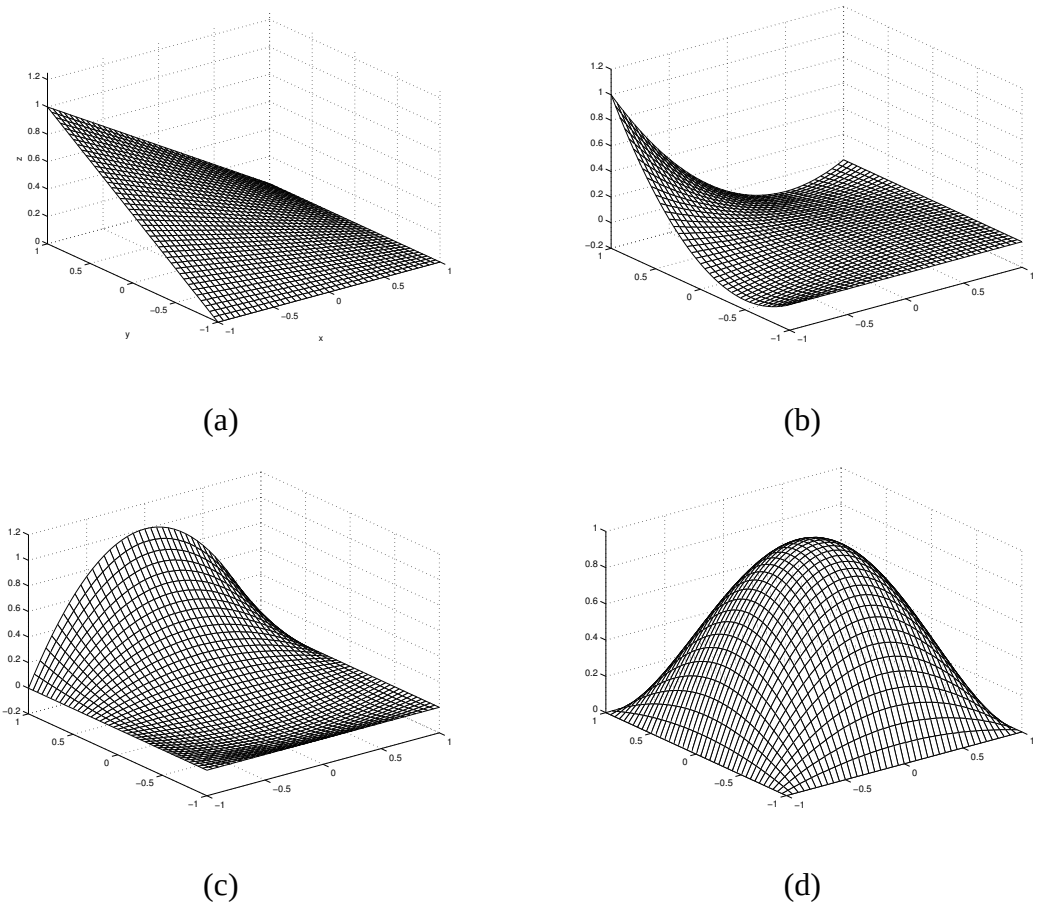

Figure 4.4: Standard finite element trial functions over the unit quadrilateral: (a) linear corner function, (b) Quadratic corner function, (c) Quadratic mid edge function, (d) Quadratic mid-element function

$f_j \in V_j$ can be found to be arbitrarily close to any function in $L_2(\mathsf{H})$. This shows the spaces $V_j, j \to \infty$, to be dense in $L_2(\mathsf{H})$, therefore satisfying condition 2. The MRA's scaling functions $\varphi_{j,k}$ that form a basis of V_j are defined as the set of bi-linear finite element functions associated with the vertices of the elements of H_j. Each function has the value 1 over its associated vertex and attenuates linearly to zero towards its immediate neighbouring vertices, as shown in figure 4.4. The relationship between the scaling function and the vertices indexed by the subdivision scheme (linear case) gives the condition $K^{'}(j) = K(j)$. Therefore, the scaling functions of the space V_j are referenced by the index set $K(j)$ defined by the linear subdivision scheme.

The Quadratic MRA on H_j

The quadratic MRA defines the spaces V_j as the set of functions that are piece-wise bi-quadratic over the elements H_j. Using the same arguments used in the linear MRA, the first two conditions of an MRA can be show to be satisfied by the functional spaces V_j. The scaling functions $\varphi_{j,k}$ that span each space V_j are defined as the finite element trial functions with a bi-quadratic representation over the elements H_j. For this, a scaling function is associated with each vertex, mid-edge and mid-face point of the faces H_j. Each function is bi-quadratic and interpolates over the vertices, mid-edge and mid-face points of the element, as illustrated in figure 4.4. The relationship between the quadratic function and the points indexed by the quadratic subdivision scheme gives the condition $K^{'}(j) = K(j)$. Therefore, as with the linear MRA, the scaling functions are referenced by the index set $K(j)$ defined by the quadratic subdivision scheme.

Some General Properties of an MRA

This sections describes some important properties of an MRA required for the development of the wavelets.

Refinement Relation:

The nested properties of V_j imply the scaling functions $\varphi_{j,k} \in V_j \subset V_{j+1}$ can be constructed from the scaling functions of V_{j+1}.

$$\varphi_{j,k} = \sum_{m \in K(j+1)} h_{j,k,m}\, \varphi_{j+1,m}. \tag{4.4}$$

Due to the interpolation properties of the scaling function used in both linear and quadratic MRAs, $\varphi_{j,k}(s_{j,l}) = \delta_{j,l}$ for $k, l \in K(j)$, the refinement relation simplifies to,

$$\varphi_{j,k} = \varphi_{j+1,k} + \sum_{m \in M(j)} h_{j,k,m} \varphi_{j+1,m}. \tag{4.5}$$

The summation now only involves the scaling functions associated with the newly created points at subdivision $j+1$.

Numerical Analysis Applications:

A function $f \in L_2(\mathsf{H})$ may be approximated with a function $f_j \in V_j$ for $j \in \{0, 1, \ldots \infty\}$. The approximation may then be represented by an expansion over the scaling functions spanning the space V_j,

$$f \rightarrow f_j = \sum_k \lambda_k \, \varphi_{j,k}, \tag{4.6}$$

where λ_k are the expansions coefficients. Increasing the value of j increases the accuracy of the approximation. As $j \rightarrow \infty$, f is fully recovered and so a basis for the the space $L_2(\mathsf{H})$ is given by,

$$L_2(\mathsf{H}) = \lim_{j \rightarrow \infty} V_j. \tag{4.7}$$

4.3.4 Linear & Quadratic Lifted Wavelets

The purpose of the wavelets are that they produce bases for spaces W_j that complement the spaces V_j in V_{j+1}. That is $V_{j+1} = V_j \oplus W_j$ for all $j \in \{0, 1 \ldots, \infty\}$. The wavelets spanning the bases W_j are formally defined by,

$$\{\psi_{j,k}, k \in M^{'}(j)\}, \tag{4.8}$$

for some index set $M^{'}(j)$.

The important property of the wavelets in this work is that the space V_j can be decomposed into the space V_0 and the spaces W_i for $i \in \{0, 1, \ldots j-1\}$.

$$V_j = V_0 \oplus_{i=0}^{j-1} W_i \tag{4.9}$$

On letting $j \rightarrow \infty$ the above summation creates a second basis for $L_2(\mathsf{H})$. The decomposition property also allows a second representation of the approximated function $f_j \in V_j$ defined in (4.6). This is given by,

$$f_j = \sum_{k \in K(0)} \lambda_{0,k} \, \varphi_{0,k} + \sum_{l=0}^{j-1} \sum_{k \in M(l)} \gamma_{l,k} \, \psi_{l,k}, \tag{4.10}$$

where $\gamma_{l,k}$ are the wavelet functions expansion coefficients.

The wavelet functions are developed in two steps. Step 1 creates a simple 'lazy wavelet' basis that spans the space W_j. These basis functions are defined as,

$$\psi_{j,m} = \varphi_{j+1,m}, \qquad \forall\, m \in M(j), \forall j \in \{0, 1, \ldots, \infty\} \tag{4.11}$$

which are the scaling functions of the space V_{j+1}, associated with the new points indexed by the set $M(j)$ defined by the subdivision scheme.

The second step modifies the wavelets in order to create a more stable Riesz basis for W_j. This process removes a contribution of the scaling functions of V_j from the lazy wavelet, a method referred to as lifting [90]. The new 'lifted wavelet' is defined by,

$$\psi_{j,m} = \varphi_{j+1,m} - \sum_{k \in K(j)} S_{j,m,k}\, \varphi_{j,k}, \qquad m \in M(j). \tag{4.12}$$

where there are no restrictions on the values of $S_{j,m,k}$ since the complementing condition between the spaces hold for arbitrary values.

Two conditions are imposed on the lifted wavelets. One requires the lifted wavelets to have zero integrals (i.e. when the functions are integrated over the surface of the hexahedron, the resulting integral is zero) as this has previously been shown to generate more stable bases for the spaces W_j [127]. The second condition requires the wavelets support to be kept to a minimum in order that they remain localised. This condition is enforced by restricting the lazy wavelet's lifting contribution to the two nearest neighbouring scaling functions of the space V_j (this is referring to each wavelet of the spaces $W_j, \forall j \in \{1, 2, \ldots, \infty\}$). Figure 4.3 presents the diagram detailing the lifting functions for the linear and quadratic lazy wavelets. These two conditions define the lifting weights of the lazy wavelets to be,

$$S_{j,k,m} = \begin{cases} \frac{\int \varphi_{j+1,m}}{2 \int \varphi_{j,k}} \text{ if } S_{j,k} \text{ is a neighbour of } S_{j+1,m} \\ 0 \text{ otherwise} \end{cases}$$

Figure 4.5 presents a typical linear lifted wavelet over a subset of H_j.

In order that these wavelet bases do not form errors in their solutions when used to discretise the angular flux, they are required to have discontinuous solutions across the planes defined by the Ω_x, Ω_y and Ω_z coordinates, see figure 4.1 for an illustration. This discontinuity is required in order to suppress the ray-effect type inaccuracies forming in 1

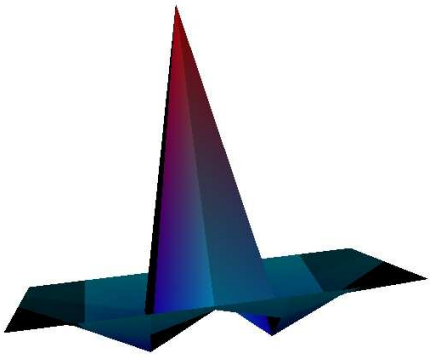

Figure 4.5: This diagram shows a typical lifted linear wavelet on a sub-domain of the hexahedron (not drawn to scale).

and 2 dimensional solutions that are caused by unstable wavelets with collocation points centred over these planes. The construction of the wavelets therefore follows the theory described in this chapter but are generated over the 8 octants of the sphere (or hexahedron) independently in order that the approximation has the desired discontinuities. This method is analogous to the method described in chapter 3 - section 3.3.5.

4.4 Numerical Examples

In order to demonstrate the hexahedral wavelets capabilities, 4 steady state two dimensional neutron transport problems are solved. The problems have been selected in order to demonstrate the angular discretisations accuracy in representing the direction of particle travel and to reveal the susceptibility to ray-effect forming in their solutions. The wavelet solutions are compared with solutions of the established methods S_N and P_N, a discontinuous Galerkin Discrete Ordinate model (DG-DO) [21] and benchmarked solutions are provided by the established even-parity code EVENT and the probabilistic model MCNP.

The following wavelet calculations have been performed with the recently developed code RADIANT [19]. The spatial variables are discretised using the finite element CVFEM

Angular Expansion sizes		
N (Angular expansion)	HLW_N	HQW_N
0	28	76
1	76	244
2	244	868
3	868	

Table 4.1: This table lists the angular expansion sizes of the hexahedral linear and quadratic wavelet discretisations.

method [88] and bare surface boundary conditions are resolved using the Riemann approach [87] described in section 2.3. In the following calculations, the linear and quadratic wavelet solutions relating to N subdivisions of the hexahedron H are denoted by HLW_N and HQW_N, respectively. The expansion sizes of the hexahedral wavelet discretisations are listed in table 4.1.

4.4.1 Problem 1: A Box Source Surrounded by a Pure Absorbing Medium

This numerical example is designed to demonstrate the hexahedral wavelets capabilities in resolving a simple 2-D neutron transport problem. The aims are compare the accuracy of the linear and quadratic wavelets and to investigate their capabilities in comparison to the conventional angular expansions.

The problem domain is presented in figure 4.6 and the corresponding material and source information is listed in table 4.2. The domain is a 3cm by 3cm square containing a material with a pure absorbing cross-section of 0.5 cm^{-1}. The central 0.5cm by 0.5cm square contains an isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and vacuum boundaries are prescribed along all edges of the problem domain. For the following calculations, the spatial dimensions were discretised using a structured quadrilateral grid consisting of 225 elements and 256 nodes.

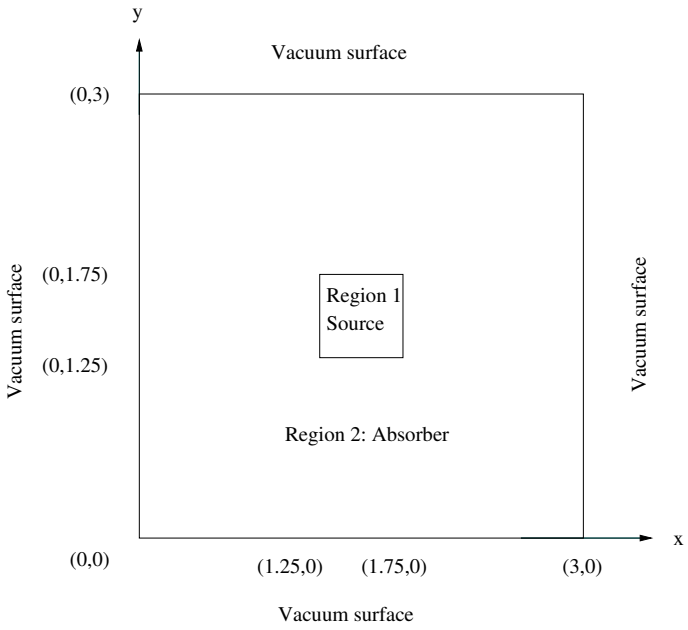

Figure 4.6: Problem 1: Diagram of the box source problem 1.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.0
2	0.0	0.5	0.0

Table 4.2: Source and material properties of Problem 1

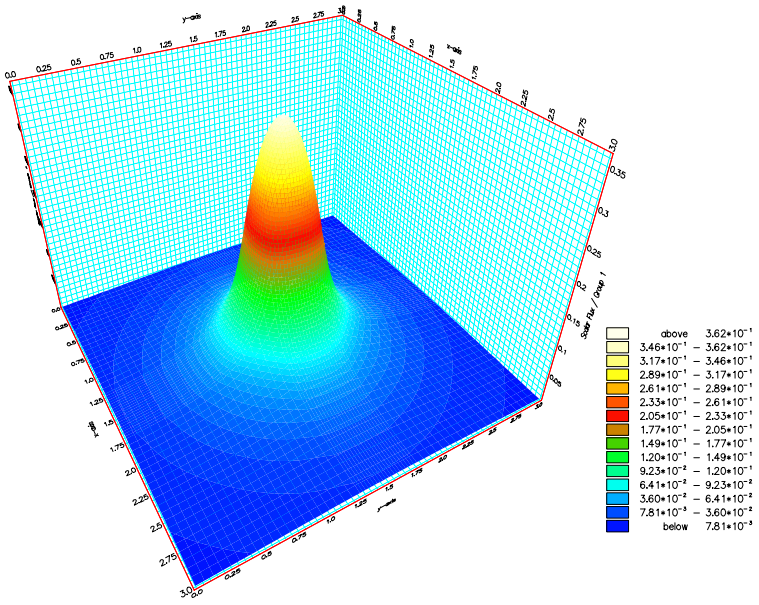

Figure 4.7: Problem 1: Exact scalar flux solution of Problem 1. This was calculated with EVENT P_{23} on a fine spatial mesh.

The benchmarked solution, presented in figure 4.7, has been provided by EVENT using a P_{21} angular approximation and a fine spatial discretisation mesh. The solution shows the scalar flux peaks at approximately 0.36 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ over the centre of the problem domain. From the centre, the scalar flux attenuates uniformly in all directions due to particles being absorbed by the material.

Figure 4.8 presents the P_5, S_6 and HLW_0 scalar flux solutions, which require 21, 24 and 28 angular basis functions, respectively. The wavelet and discreet ordinate solutions peak in the problem's centre at 0.37 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, which is in close agreement with the benchmarked solution. However, the maximum flux using P_5 falls far short of this value at approximately 0.26 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This shows that although the angular expansions are of similar size, the wavelets and S_N have resolved the source region far more accurately than P_N. This is exemplified by the graph presented in figure 4.9. This graph plots the scalar flux through the points (0,0) and (3,3) and reveals P_5 inability to correctly resolve the source region. The percentage errors using P_5, S_6 and HLW_0 at the solutions centre are 25.7%, 2.7% and 2.7%, respectively.

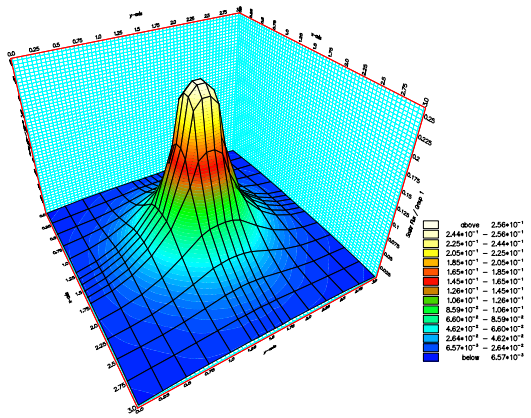

(a) P_5 RADIANT

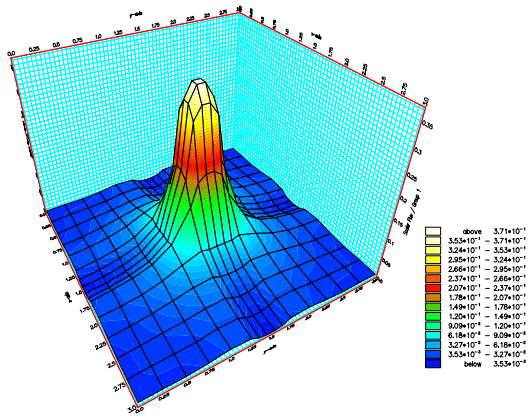

(b) S_6 RADIANT

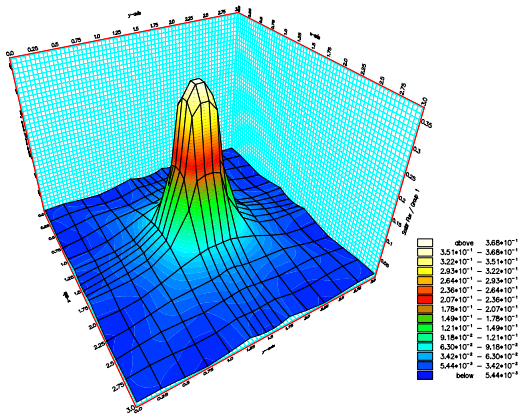

(c) HLW_0 RADIANT

Figure 4.8: Problem 1: The scalar flux solutions of the (a) P_5, (b) S_6, and (c) HLW_0 discretisations.

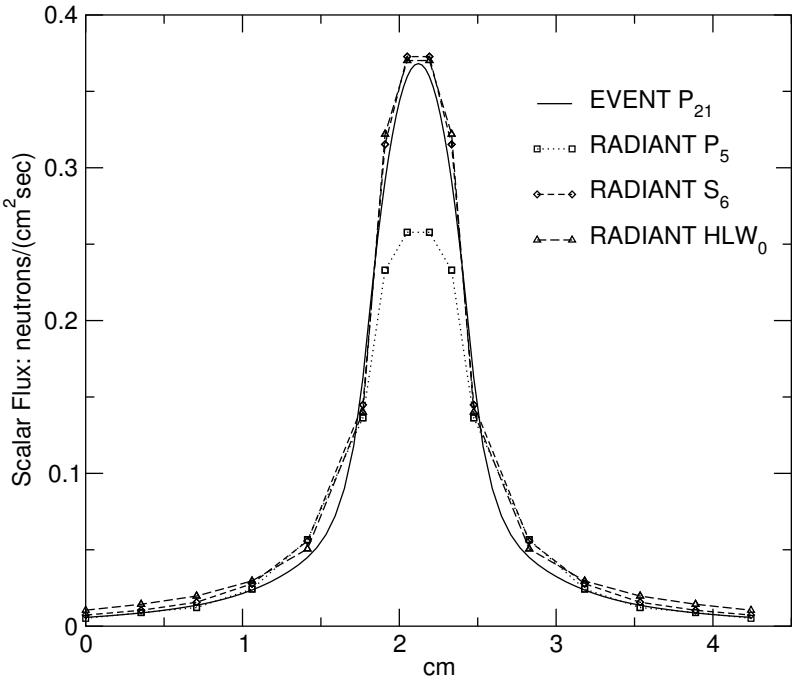

Figure 4.9: Problem 1: The P_5 S_6 and HLW_0 scalar flux solutions through the points (0,0),(3,3).

Despite resolving the source region accurately, the wavelets and discrete ordinate discretisations contain un-physical oscillations (or ray-effects) in their scalar flux solutions. These errors result from the discretisations not being rotationally invariant, the condition required in order that solutions are free from ray effects. In order to determine the full extent of the solution's ray effect distortions, figure 4.10 presents the contours of the HLW_0, HLW_1, HQW_0, S_6, S_{10} and the benchmarked scalar flux solutions.

The HLW_0 and S_6 discretisations develop large oscillations that particularly dominate the outer regions of the problem's solution. However, the oscillations can be seen to develop much closer to the source region in the S_6 solution, indicating the discrete ordinate discretisation to be more prone to developing ray effects. The HLW_0 contours just outside the source region also appears to be more accurately resolved than those of S_N.

The HLW_1 and S_{10} contours show significant improvements in their solutions over the lower order approximations. The HLW_1 discretisation display no visible signs of ray effects in its solution. However, the S_{10} discretisation continues to display a number of small oscillations, particularly in the horizontal and vertical directions from the source. The HLW_1 discretisation also agrees more closely with the benchmarked solution than S_{10} and has managed to develop the diamond shaped contours outside the source region. The HQW_0 discretisation yields similar results to HLW_1 for which the contours also reveal few visible signs of ray effects in its solution.

4.4.2 Problem 2: A Box Source Surrounded by Pure Absorbing and Scattering Materials

This numerical example is designed to demonstrate the wavelets capabilities (including investigating their susceptibility to forming ray effects) in resolving problems containing scattering cross-sections. The default linear and quadratic wavelets are compared against the conventional P_N and S_N solutions.

The problem domain is presented in figure 4.11 and the corresponding source and material information is listed in table 4.3. The problem is a 4cm by 4cm square. The central

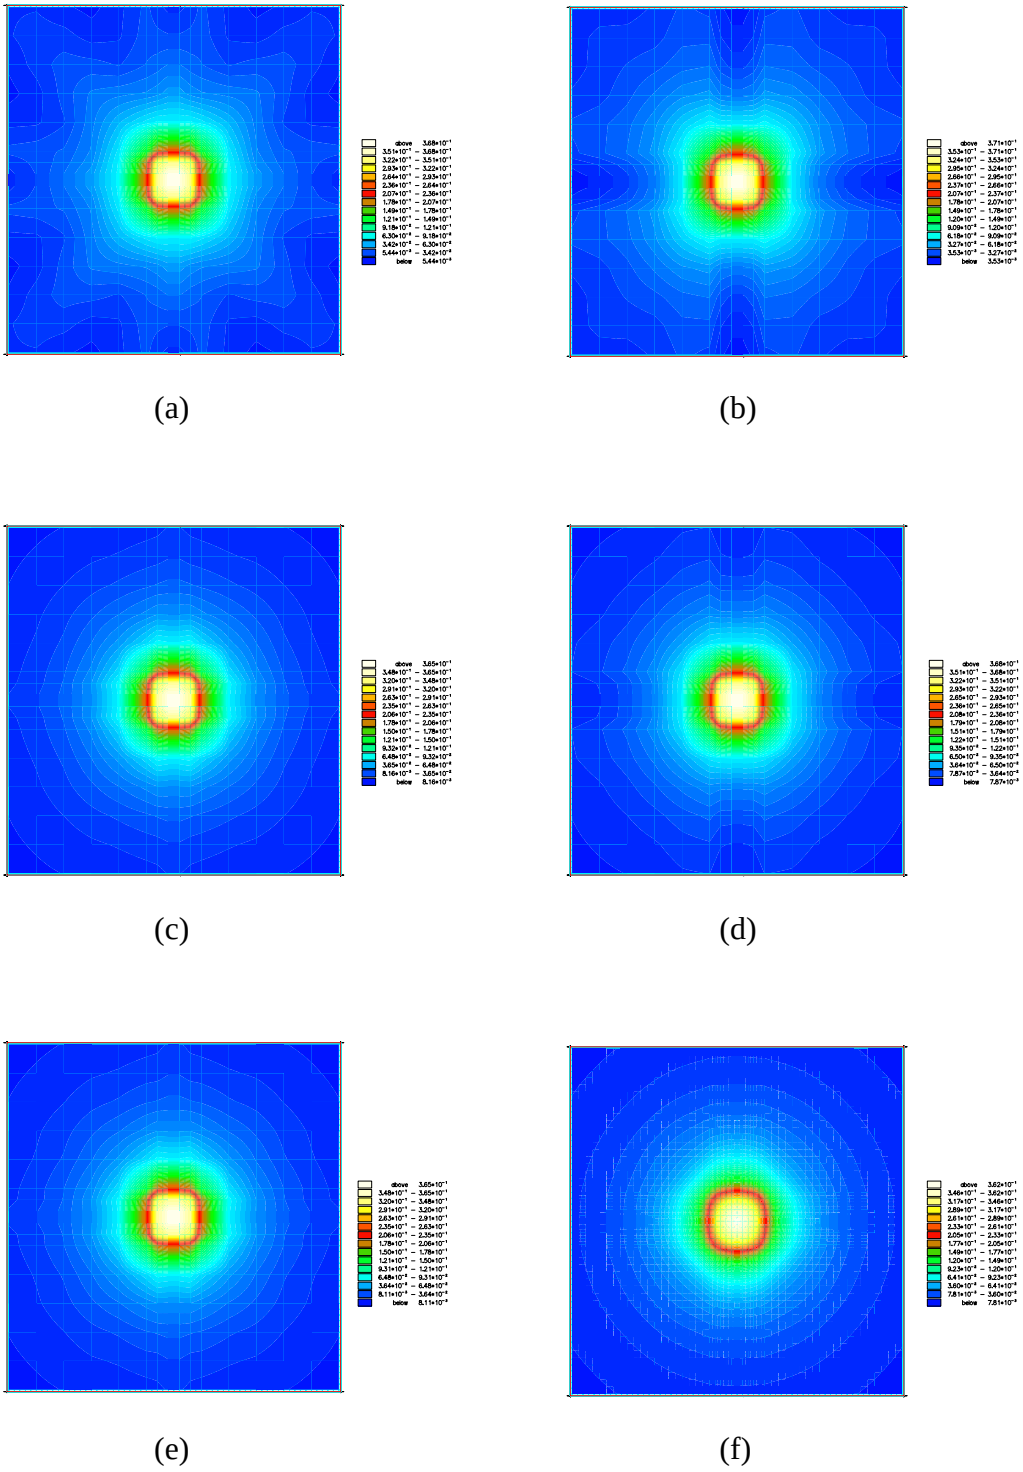

(a) (b)

(c) (d)

(e) (f)

Figure 4.10: Problem 1: The contours of the scalar flux solutions: (a) HLW_0 (28), (b) S_6 (24), (c) HLW_1 (76), (d) S_{10} (60), (e) HQW_0 (76), (f) benchmarked solution. The (.) denotes the number basis functions used in the angular expansion.

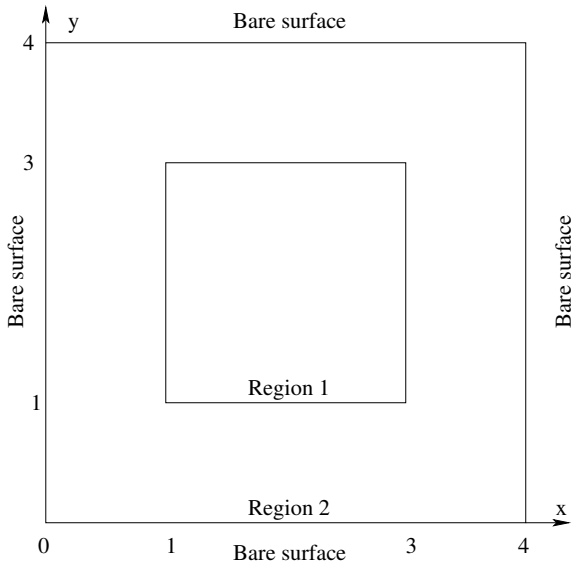

Figure 4.11: Diagram of problem 2.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	1.0 / 0.6	0.0 / 0.4
2	0.0	1.0 / 0.6	0.0 / 0.4

Table 4.3: Source and material properties of Problem 2

2cm by 2cm square contains an isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The problem's material has a total cross-section $\sigma_t = 1.0\ \mathrm{cm}^{-1}$ and a scattering cross-section which is varied throughout testing. Vacuum boundary conditions are prescribed across all edges and the spatial domain is discretised using a finite element mesh consisting of 225 elements.

The analysis of this numerical example concentrates on the wavelets ability to resolve the flux leakage through the top boundary. This is shown as line 1 in the diagram presented in figure 4.11. Comparisons of the wavelet solutions are made with those of the established S_N discretisation. A Benchmarked solution is provided by EVENT using a P_{23} angular approximation and a fine spatial discretisation mesh.

Figure 4.12 presents the benchmarked solution's flux leakage through the top boundary

for a non scattering problem. The profile shows a Gaussian shaped curve peaking at 0.043 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ over the central region. Included in this figure are the approximations from the HLW_0, HLW_1, HQW_0, S_6 and S_{10} discretisations.

The lower order HLW_0 and S_6 solutions are shown to have succeeded in capturing the general shape of the benchmark solution. Their peak fluxes over the centre of the boundary agree closely with the true value, that is, the HLW_0 and S_6 peak at 0.044 and 0.042 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, respectively. Across the boundary their profiles follow the general profile of the benchmark solution. However, both approximations have formed a number of large unphysical oscillations caused by ray effects. The graph shows the oscillations to be more severe for S_6, implying again that the wavelets solutions are less susceptible to ray effects forming in their solutions. Over the interval considered on this problem the HLW_1 is generally more accurate than S_N.

The graphs presenting the HLW_1, HQW_0 and S_{10} solutions are shown to be in close agreement with the benchmark solution. The peaks are almost exact for the three approximations. However, the wavelets have performed better in resolving the outer regions of the boundary. The scalar solutions have all formed ray effect oscillations, although they are less profound than those in the low order approximations. The solutions experience their largest oscillations over the positions $x = 1$ and $x = 3$. The S_N forms the largest variations, this once again indicates the S_N discretisation to be more susceptibility to forming ray effects in its solutions.

Figure 4.13 presents the HLW_0, HLW_1, HQW_0, S_6 and S_{10} solutions when the problem had a material scattering cross-section $\sigma_s = 0.4$. The benchmarked solution that is included in the graphs show another Gaussian shaped curve which peaks over the boundary's centre at 0.082 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

The HLW_0 and S_6 discretisations have again formed the general features of the benchmark solution for which their scalar flux peaks at 0.083 and 0.081 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, respectively. Both solutions contain ray effect oscillations, however, these are not as severe as the oscillations that formed in the non-scattering calculations. This is mainly due to the higher flux isotropy that is induced by the scattering cross-section, which in turn present an simpler flux to approximate. The graph also shows the oscillations to be more

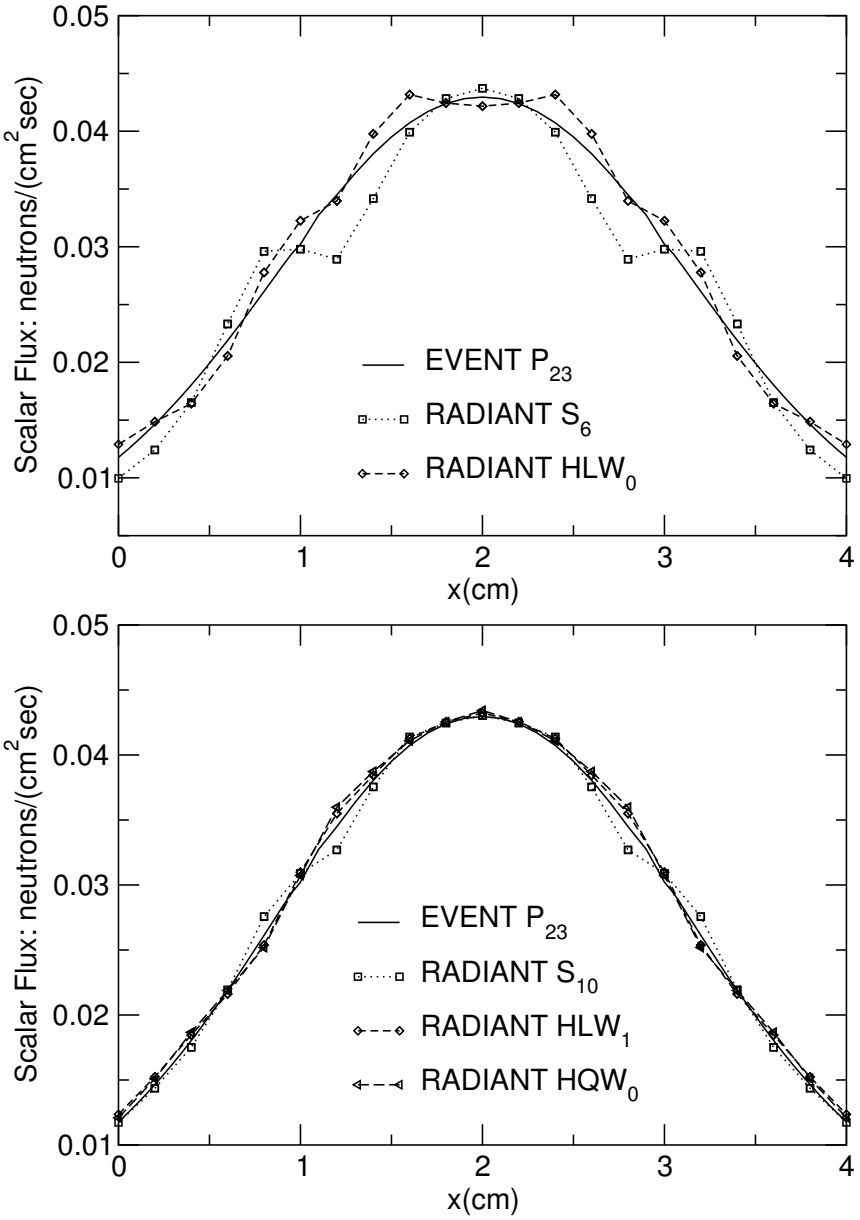

Figure 4.12: Problem 2: The scalar flux leakage through the top boundary of problem 2 with a scattering cross-section σ_s = 0.0. Top graph: S_6 and HLW_0 solutions using 24 and 28 angular basis functions respectively. Bottom graph: S_{10}, HLW_1 and HQW_0 solutions using 60, 76 and 76 angular basis functions respectively.

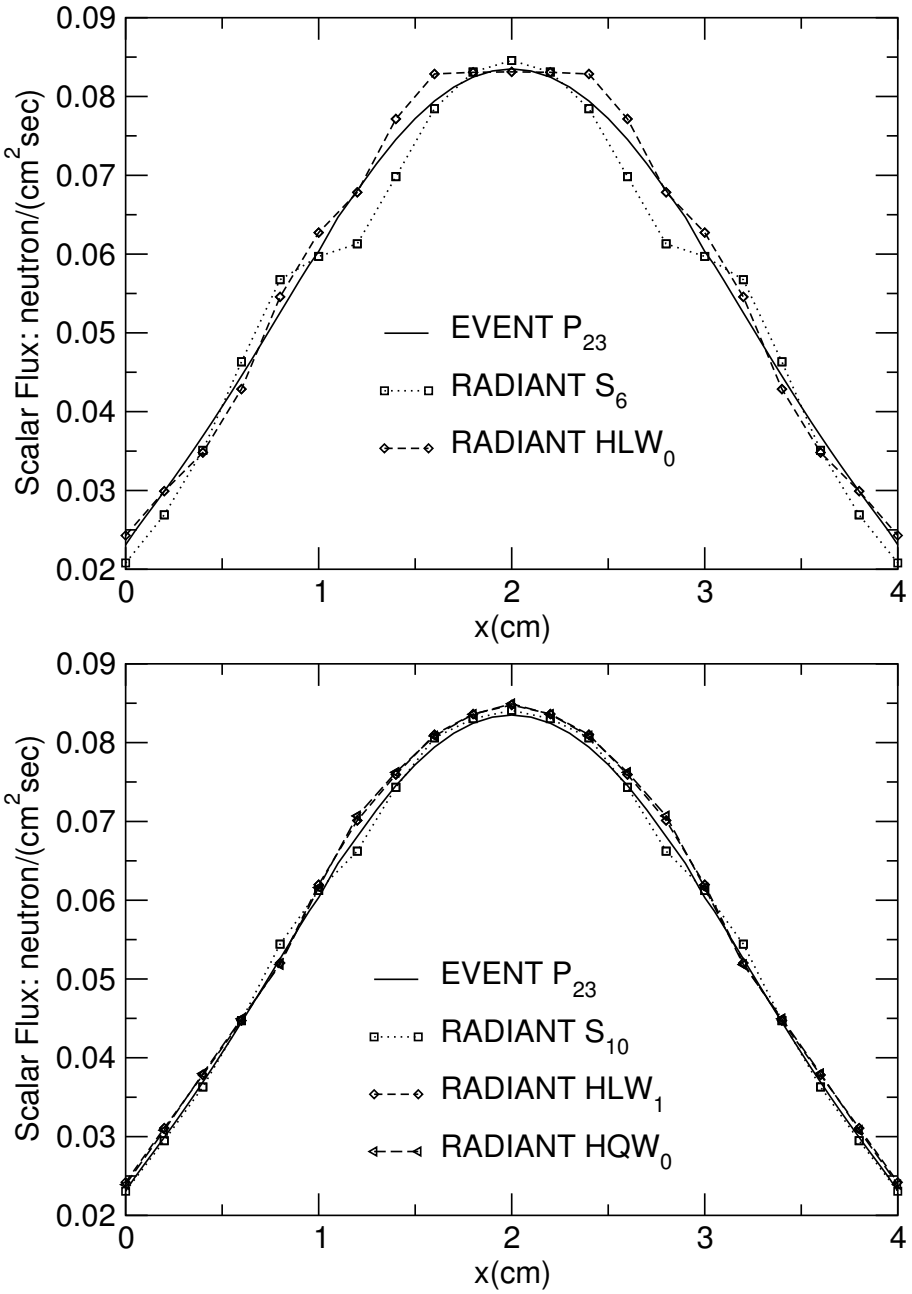

Figure 4.13: Problem 2: The scalar flux leakage through the top boundary of problem 2 with a scattering cross-section σ_s = 0.4. Top graph: S_6 and HLW_0 solutions using 24 and 28 angular basis functions respectively. Bottom graph: S_{10}, HLW_1 and HQW_0 solutions using 60, 76 and 76 angular basis functions respectively.

dominant in the S_6 solution, continuing the trend that the discrete ordinate method is the most susceptible to ray effects. The bottom graph of figure 4.13 shows close agreements between HLW_1, HQW_0, S_{10} and the benchmark solution. Small oscillations caused by ray effects have formed in all flux profiles about the points $x = 1$ and $x = 3$. These again are more predominate in the S_{10} approximation.

Figure 4.14 presents the flux leakage error for the HLW_0, HLW_1 and HQW_0 discretisations when the problem's scattering cross-sections were $\sigma_s = 0.0$ and $\sigma_s = 0.4$. The HLW_0 maximum error of 9.0 $\%$ for the non-scattering problem can be reduced to 5 $\%$ upon increasing the angular expansion to HLW_1 or increasing the wavelet polynomial representation to quadratic variation HQW_0. The scattering calculations show similar results where the maximum error for HLW_0 (6$\%$) can be reduced to 5$\%$ or 4$\%$ using the HLW_1 and HQW_0 discretisations, respectively.

4.4.3 Problem 3: Straight Ducts Surrounded by Pure Absorbing Materials

This problem is designed to demonstrate the capabilities of the wavelet expansions to accurately resolve problems with more demanding, non uniform domains containing void material cross-sections. In this example the linear wavelets are considered and solutions are compared to those of the established methods - EVENT and RADIANT P_N.

The problem domain is presented in figure 4.15, and the source and material information is listed in table 4.4. The problem has a rectangular domain of size 36cm $\times$ 28cm. The central 6cm $\times$ 6cm square region contains an isotropic source of strength 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and a purely absorbing material with a cross-section $\sigma_t = 0.5$. Extending from the source are two straight ducts containing voids. Both ducts have a width of 6cm that run to the upper and lower boundaries. The material surrounding the ducts have a pure absorbing cross-section $\sigma_t = 0.5$. Vacuum boundaries are prescribed along all edges and the spatial variables are discretised using a mesh of 900 finite elements.

The problem domain was solved using HLW_0, HLW_1, HQW_0, P_5, P_9 and P_{11}. Their

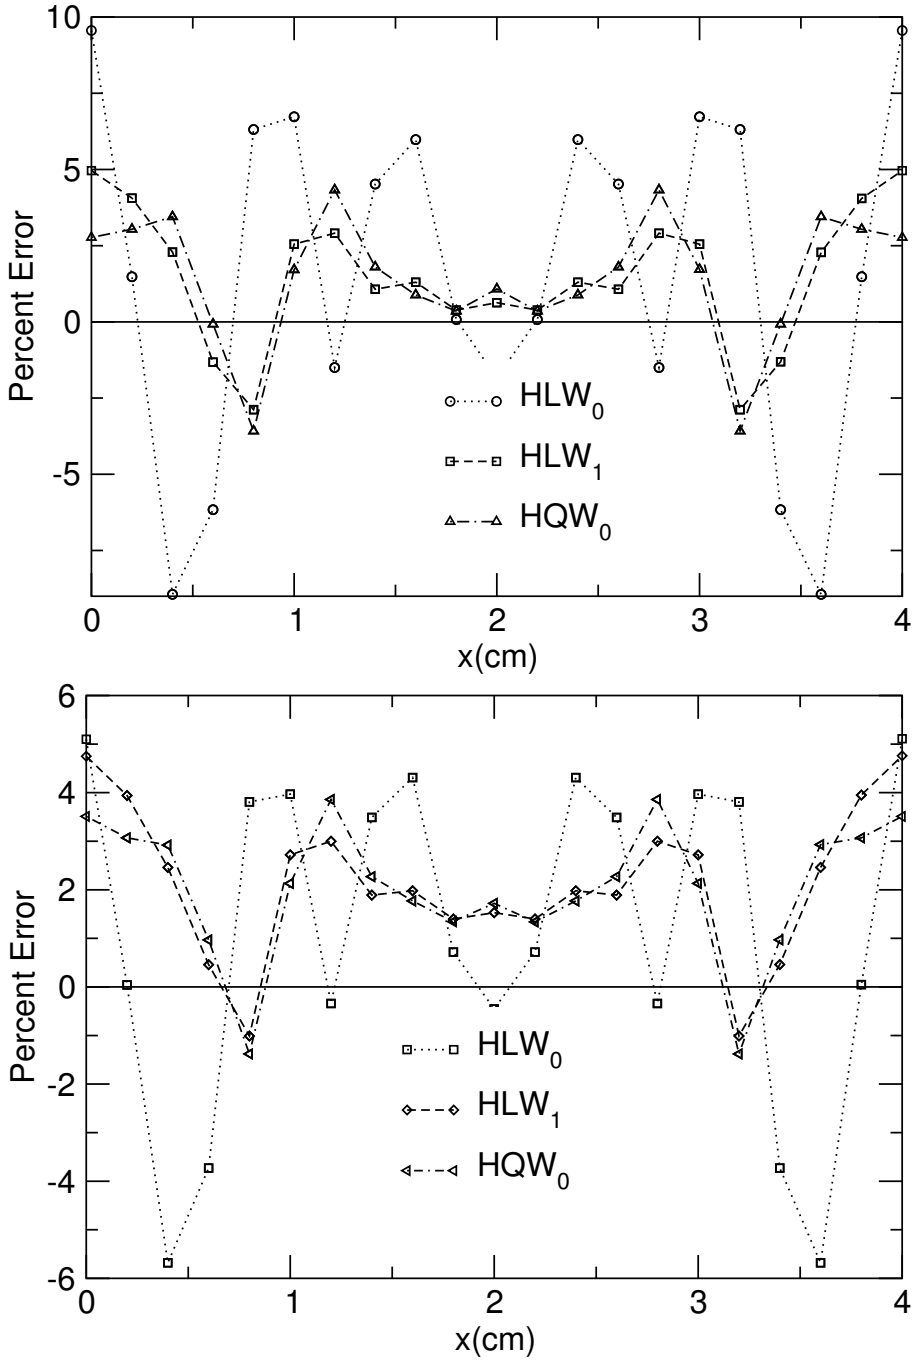

Figure 4.14: Problem 2: The error graphs of the HLW_0, HLW_1 and HQW_0 solutions of Problem 2. The top graph presents the errors for $\sigma_s = 0.0$. The bottom graph presents the errors for $\sigma_s = 0.4$.

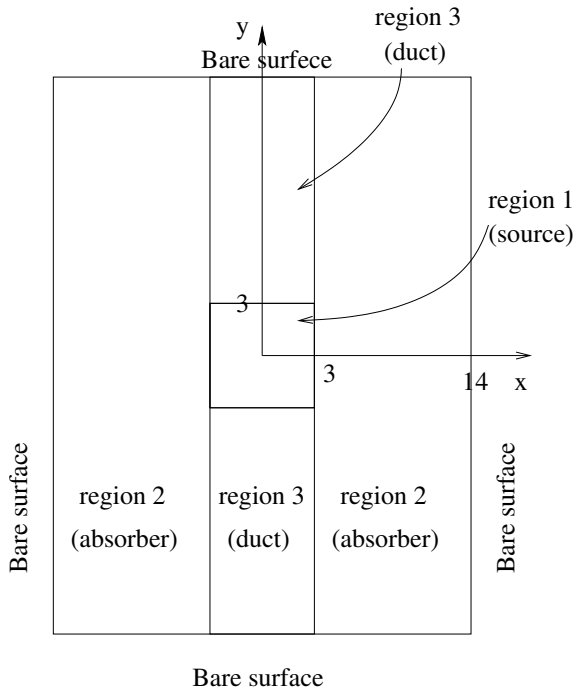

Figure 4.15: Problem 3: Diagram of the straight ducts surrounded by a pure absorbing material problem.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.0
2	0.0	0.5	0.0
3	0.0	0.0	0.0

Table 4.4: Source and material properties of Problem 3: The Straight Duct Problem.

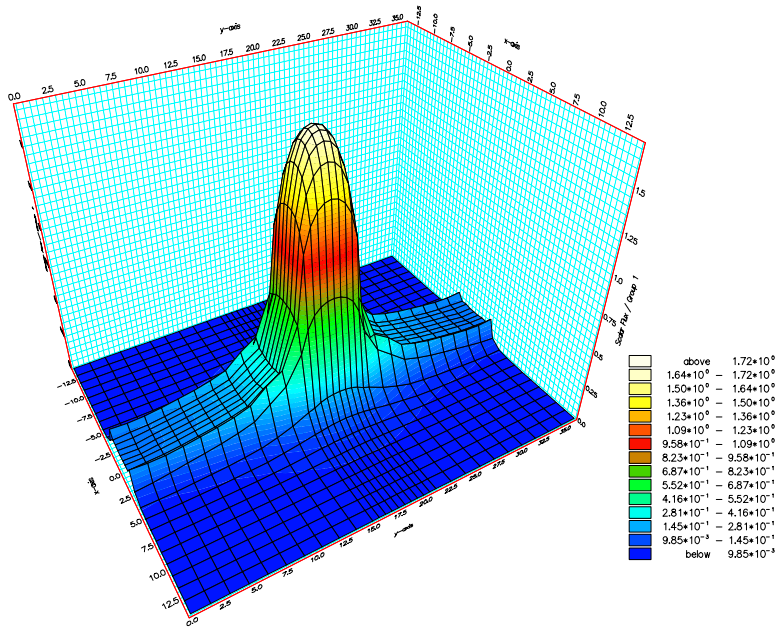

Figure 4.16: Problem 3: Scalar flux profile of the P_5 RADIANT solution.

scalar flux solutions are presented in figures 4.16-4.20. The profiles, which are all similar, have peaks in the scalar flux forming over the source region at approximately 1.7 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. From the source region there is a high attenuation of flux where particles stream through the absorbing materials. Through the ducts the flux declines at a slower rate as the particles propagate down them without collision.

The wavelet solutions have formed a number of un-physical oscillations through the ducts. However, these are not necessarily angular discretisation errors as some oscillations are also present in the P_N solutions. One example is the raised ridges running the length of the border between the ducts and absorbing regions. These oscillations, which form in all solutions, are most likely a result of Gibbs oscillations that form through the spatial discretisation. This is illustrated in figure 4.21 which presents the P_{11}, HLW_1 and HQW_0 flux leakage through the top boundary. The P_{11} solution reveals the extent of the Gibbs oscillations (induced by the spatial discretisation) corrupting the scalar flux profile. Not only are there raised ridges running the length of the border between the duct and absorbing material but there then follows a series of oscillations within the duct region. This includes a sharp drop from the raised ridge from 1.6 to 0.9 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The wavelets solutions have also displayed similar Gibbs oscillations. However the initial

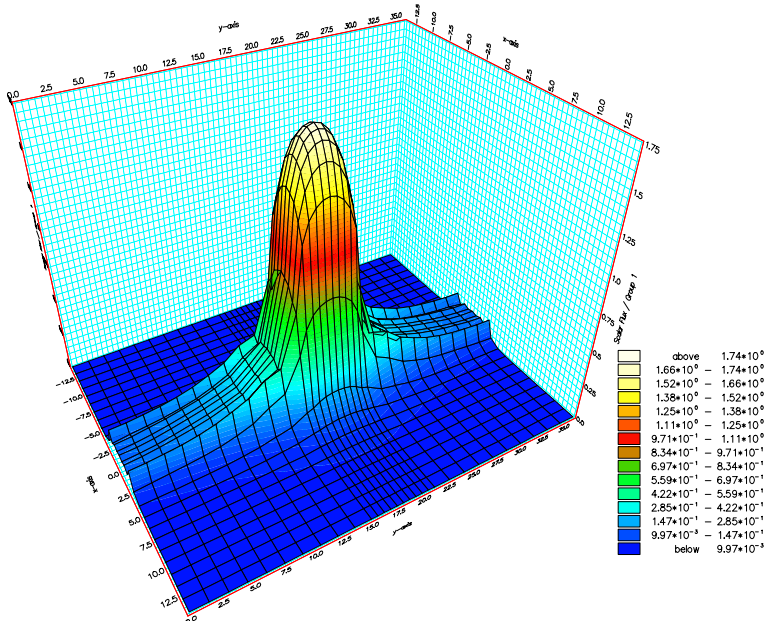

Figure 4.17: Problem 3: Scalar flux profile of the P_{11} RADIANT solution.

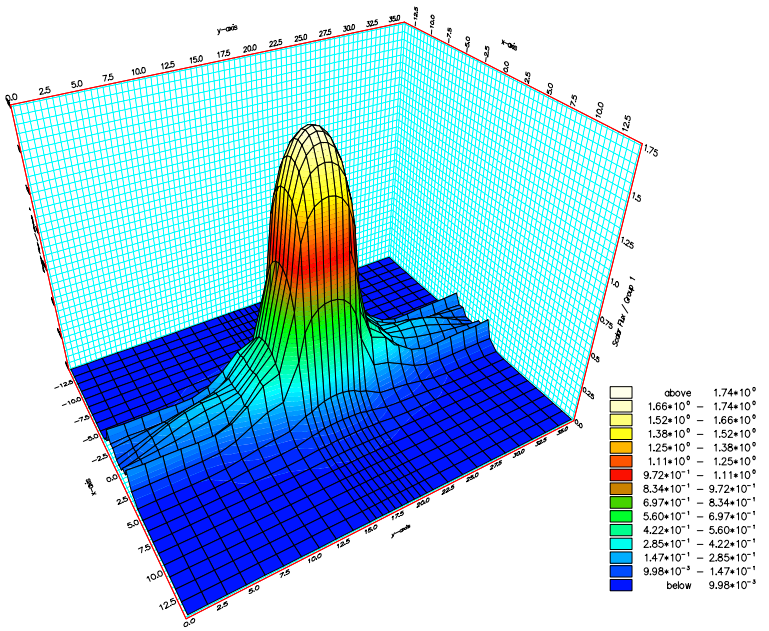

Figure 4.18: Problem 3: Scalar flux profile of the HLW_0 RADIANT solution

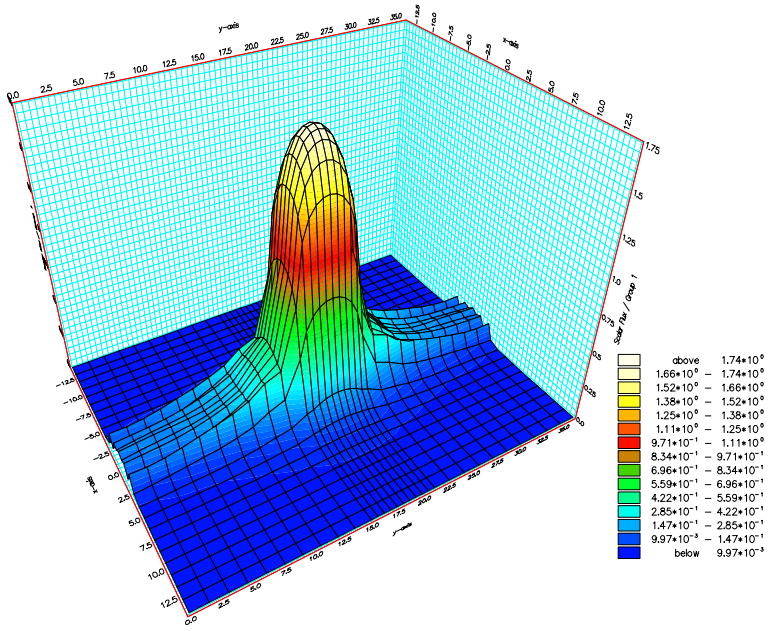

Figure 4.19: Problem 3: Scalar flux profile of the HLW_1 RADIANT solution

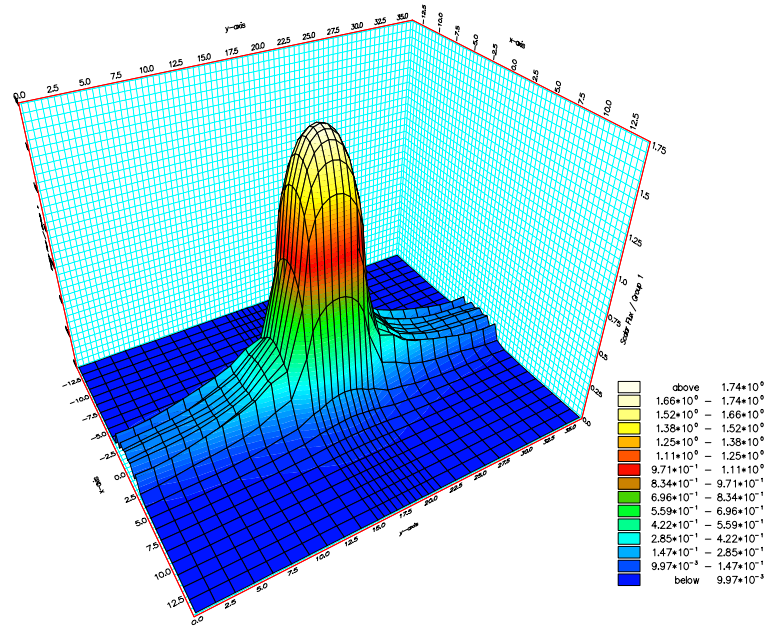

Figure 4.20: Problem 3: Scalar flux profile of the HQW_0 RADIANT solution

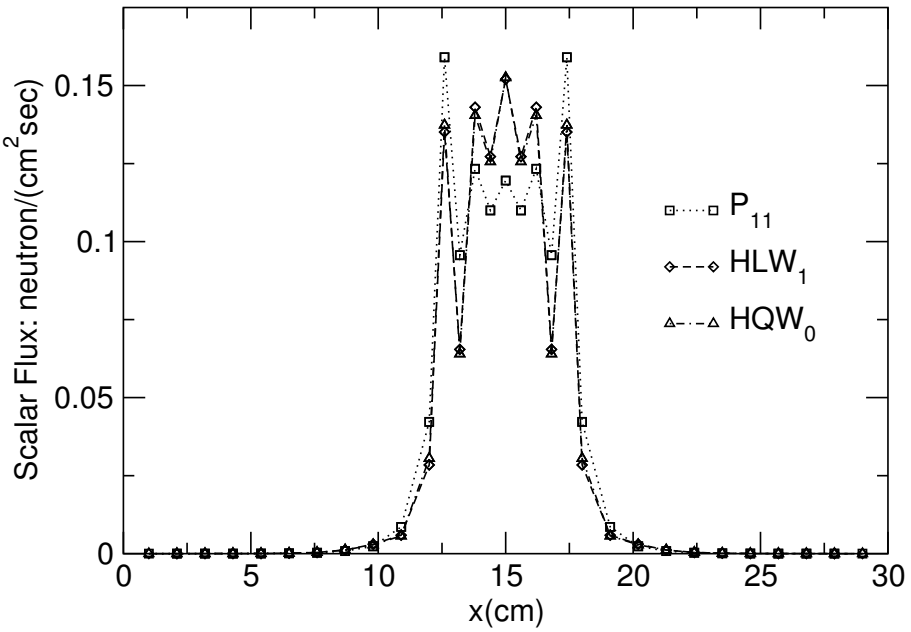

Figure 4.21: Problem 3: Scalar flux profile through the line (-16,18),(16,18) of the hexahedral wavelet and P_N solutions.

peak and trough at the border between the duct and absorbing material are lower than that found with P_N, reaching 1.4 and 0.6 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, respectively.

In order to investigate the other oscillations, figure 4.22 presents the graphs plotting the scalar flux from position (0,0) to position (0,18). This plots the flux from the problem's centre, through the top duct to the top boundary. Included in the graph are the DG-DO solution using 2,500 directions.

The sharp attenuation in scalar flux over the interval (2,7.5) reveals the largest variations between the wavelets and P_N. The wavelet solutions appear to be close to convergence over this region since increasing the approximation has little effect on the scalar flux. This is confirmed by the error graph presented in figure 4.23. The graph shows the wavelet solutions to have the smallest error over this region in comparison to P_N, and that increasing the expansion has little effect on the error. This indicated that the wavelet solution is close to angular convergence on the given spatial grid.

Over the interval (7.5,18) the HLW_0 solution forms large oscillations in its scalar flux.

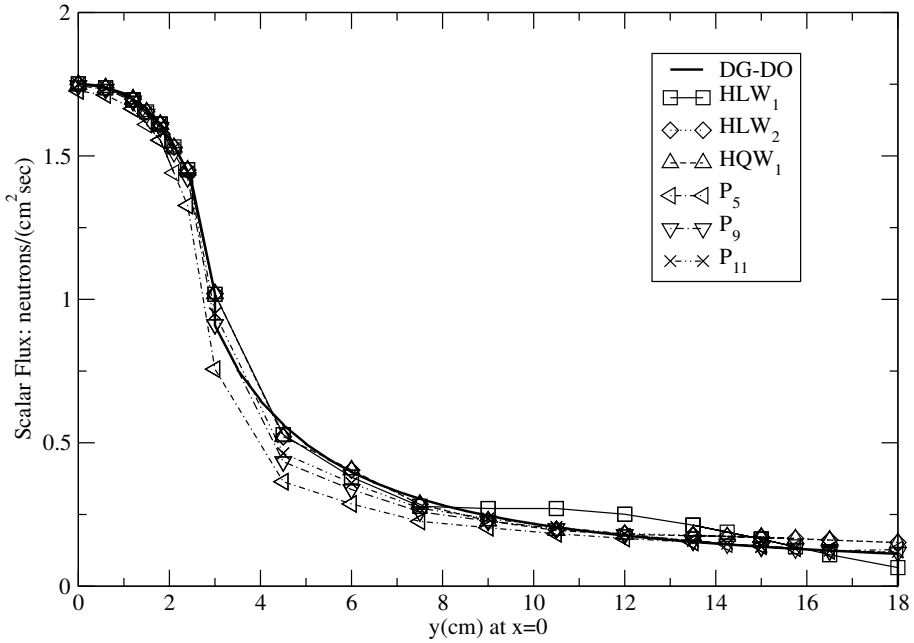

Figure 4.22: Problem 3: Scalar flux profile through the line (0,0),(0,18) of the hexahedral wavelet and P_N solutions.

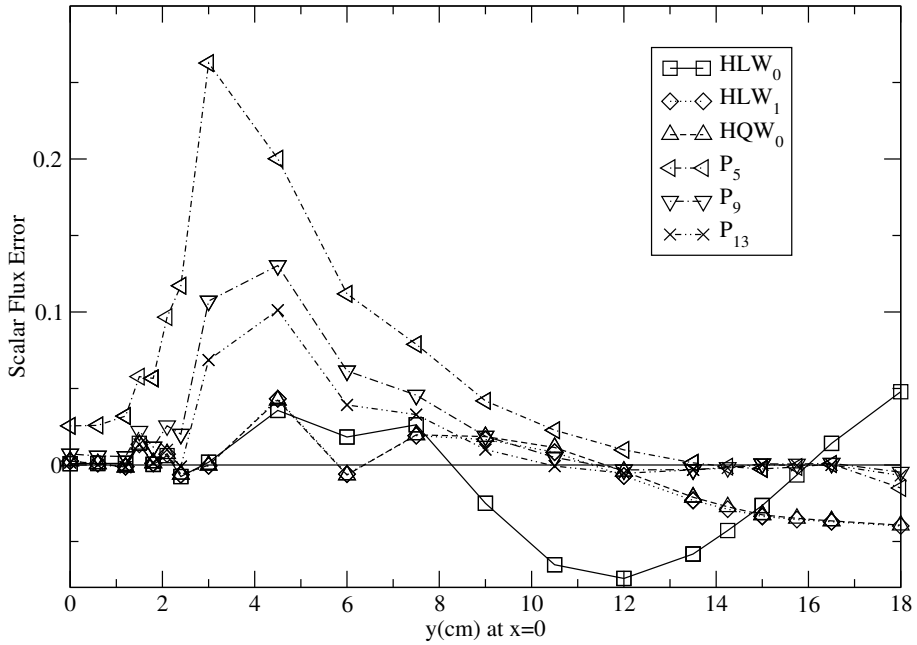

Figure 4.23: Problem 3: Scalar flux error profile through the line (0,0),(0,18) of the hexahedral wavelet and P_N solutions.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.0	0.0
2	0.0	0.0	0.0

Table 4.5: Source and material properties of Problem 4: The cylinder source inside a cylinder.

These fluctuations are mitigated by increasing the approximation to HLW_1 or HQW_0, indicating that the oscillating errors are a fault of the angular discretisation. This again agrees with the difference plot in figure 4.23. This graph also shows that the linear and quadratic wavelets produce very similar solutions as the errors are almost identical over the considered range of the problem. However, the P_N solutions are in close agreement with each other over this range, showing convergence has almost been attained with this discretisation. For P_N, a scalar flux of 0.125 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ forms the boundary point $x = 18$, this is just short of the wavelet solution of 0.15 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

4.4.4 Problem 4: A cylinder Source within a Cylinder

This Problem is designed to demonstrate the wavelets capabilities in solving demanding geometries using unstructured spatial meshes. In this demonstration the linear wavelets are compared with the solutions of EVENT and MCNP.

The problem domain is presented in figure 4.24 and the corresponding source and material information is listed in table 4.5. The problem domain is an infinitely long cylinder of radius 15cm. The central cylinder of radius 1cm contains an isotropic source of intensity 1.0 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The whole domain is a void, that is $\sigma_t = 0.0$, and a vacuum boundary condition is applied across the external surface. For the following calculations the spatial variables are discretised using Delauney triangulation [131]. The spatial mesh, presented in figure 4.25, consists of 1350 triangular elements using 676 nodes.

This problem was analysed in the numerical examples of chapter 3 and revealed some unexpected results. It was found that P_N was unable to resolve this problem due to its low

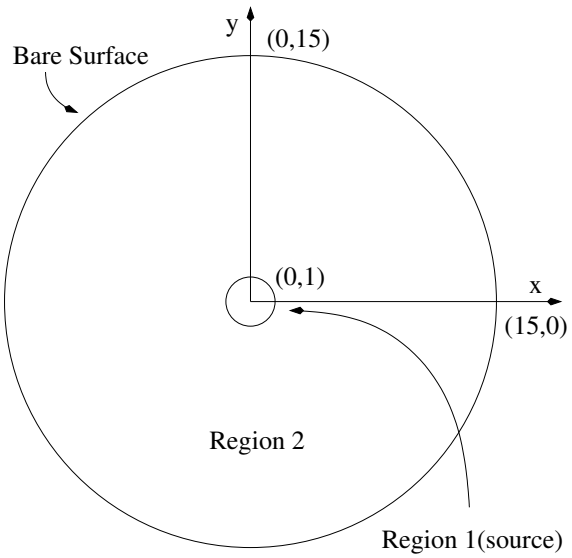

Figure 4.24: Problem 4: Diagram of the cylinder source within a cylinder problem.

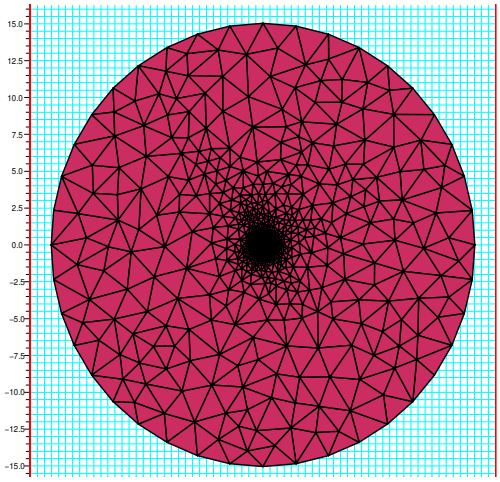

Figure 4.25: Problem 4: Diagram of the spatial mesh used in the wavelet and S_N calculations of problem 4. The mesh was generated by means Delaunay triangulation using 1250 element and 676 nodes.

cross section. In fact not even a P_{19} discretisation could provide an approximation that was close to the exact solution. Extensive tests were performed in order to provide enough convincing evidence that the S_N discretisation could resolve the domain accurately. On the basis of the previous chapter's findings it is assumed here that the S_N approximations are converging to the correct solution. The MCNP calculations have also been included in order to give further verification to these solutions.

The problem was solved using HLW_0, HLW_1, HQW_0, S_6 and S_{10}. Figure 4.26 presents the S_6 and S_{10} scalar flux solutions. These profiles form a sharp maximum in the scalar flux over the source region reaching approximately 1.5 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. From this peak, the flux undergoes a steep decent over the boundary of the source and void regions. Once outside the source region the scalar flux either drop to zero or propagates along the S_N directions. The wavelet approximations, presented in figure 4.27, show similar profiles to S_N. The peaks in scalar flux climb to approximately 1.5 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ over the centre of the domain which then falls sharply over the outer regions of the source. Once leaving the source region, the flux's sharp gradient levels off and follows a gentle decline towards the problem's boundary.

The graph presented in figure 4.28 compares these scalar flux profiles with MCNP plotted cross the line from points (0,0) to (0,15). The P_9 solution has also been included in order to highlight the method troubles in resolving this numerical example. The graph shows the S_N and wavelet solutions to in close agreement over the source region with their peak fluxes just exceeding 1.55 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This value is in close agreement with MCNP which registered a peak flux of 1.57 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Across the edge of the source region, the S_N, wavelets and MCNP solutions form a sharp decline in scalar flux. These solutions then level off into a gradual gradient towards the problem's boundary. The P_9 solution shows vast differences in its scalar flux profile and has no common features with the other solutions (this solution was shown in [132] to be highly inaccurate). The profile peaks far below the value predicted by MCNP, reaching just 0.33 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The solution also fails to display a steep gradient across the source edge. In fact, the P_9 solution shows the flux to take a smooth decent towards the problem's boundary.

The wavelet and S_N discretisations have formed a number of ray effects in their solu-

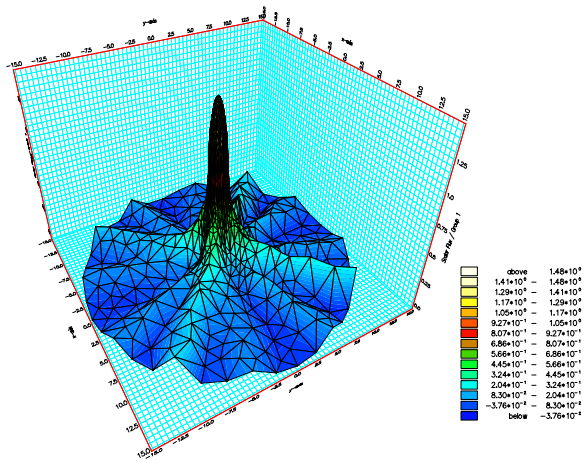

(a) S_6 solution

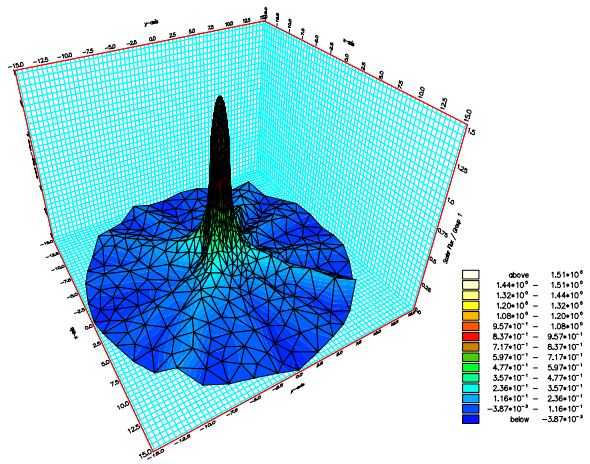

(b) S_{10} solution

Figure 4.26: The (a) S_6 and (b) S_{10} scalar flux solutions of problem 4.

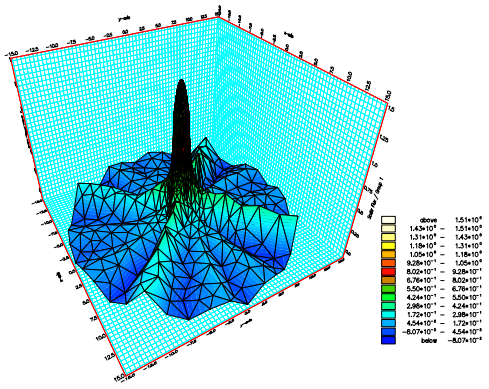

(a) HLW_0 solution

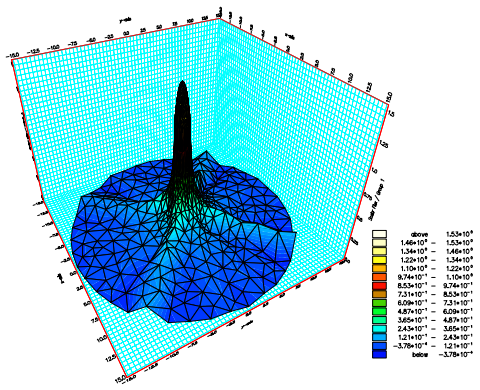

(b) HLW_1 solution

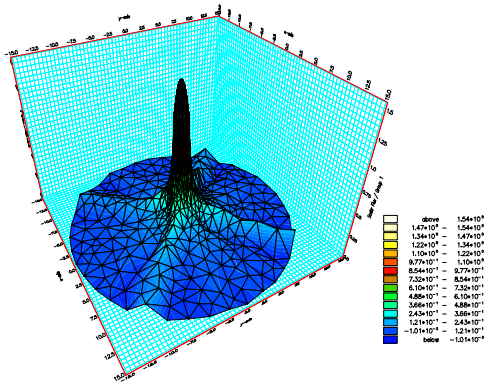

(c) HQW_0 solution

Figure 4.27: The (a) HLW_0, (b) HLW_1 and (c) HQW_0 scalar flux solutions of problem 4.

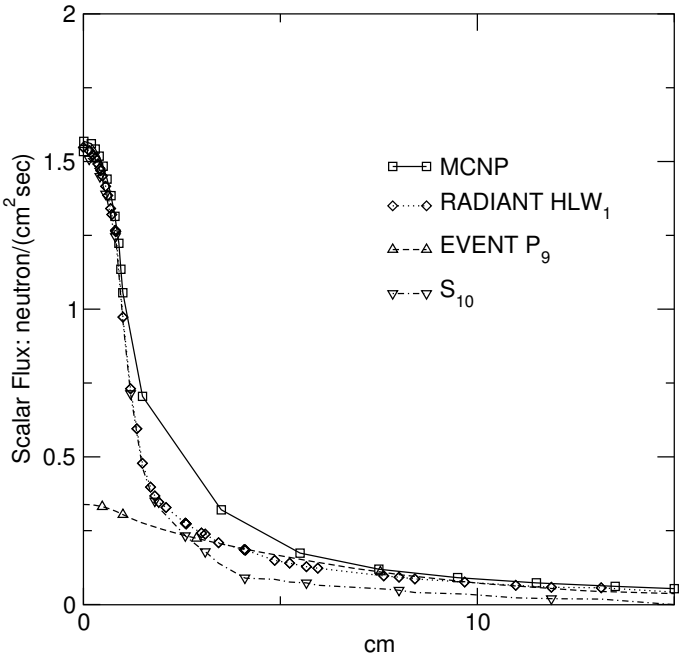

Figure 4.28: Problem 4: The scalar flux profiles of the EVENT P_9, RADIANT HLW_1 and MCNP solutions. The profile was taken along the line crossing through the centre of the problem domain.

tions of this numerical example. Figure 4.29 presents their solution's flux contours which clearly depict ray effect distortions as fingers of flux extending from the source region. The figure shows both the S_6 and HLW_0 discretisation have performed poorly in resolving some regions of this problem and that the oscillations dominate outer regions of the domain. In fact, both solutions have large sections of the outer domain that do not experience flux. However, it appears that this has occurred more often for S_6 than HLW_0.

The HLW_1, HQW_0 and S_{10} solutions have also developed ray effects. However, it again appears that S_{10} is more prone to these errors than the wavelets. It can be seen that both wavelet solutions cover more of the outer regions of the domain with flux. This again provides further evidence that the S_N discretisation suffers more to ray effects than the wavelets.

4.5 Conclusions

In this chapter two new spherical wavelet bases for discretising the angular variable of the Boltzmann transport equation has been presented. The wavelet bases are based on the ideas of Sweldens second generation wavelets and approximate the angular flux over the surface of the Hexahedron. The spherical wavelet's capabilities for accurately representing the angular dependence of the angular flux have been demonstrated by solving four 2-D numerical examples. The numerical examples were selected in order to demonstrate the wavelets abilities to resolve isotropic sources, scattering cross section as well as particle streaming through voids.

The wavelet solutions were compared with the solutions of the established S_N and P_N methods (where the same code discretising the spatial domain was used to ensure fair comparisons). Other established radiation transport codes were also used to provide benchmarked solutions.

The numerical examples have demonstrated that the wavelet bases have the potential to become an established method in neutron transport approximations. They constantly performed well in all numerical examples and were shown to be as accurate as the S_N and

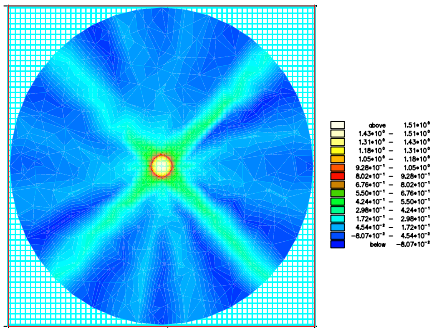

(a) HLW_0 solution

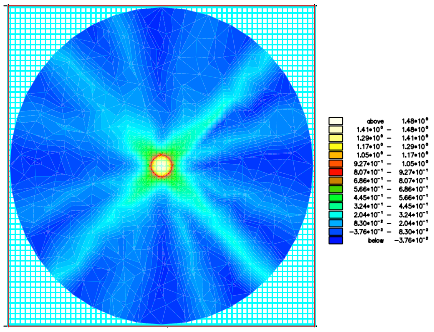

(b) S_6 solution

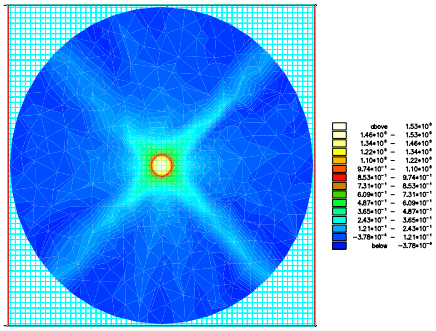

(c) HLW_1 solution

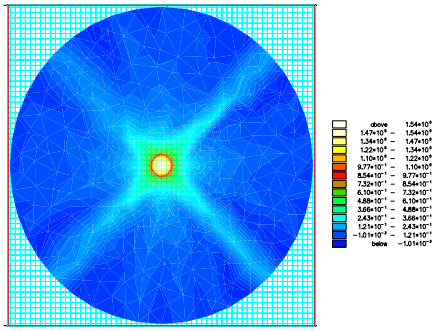

(d) HQW_0 solution

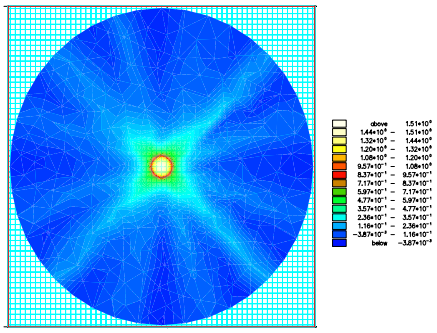

(e) S_{10} solution

Figure 4.29: Problem 4: Scalar flux contours of the solutions to problem 4. (a) HLW_0 (28), (b) S_6 (24), (c) HLW_1 (76), (d) HQW_0 (76), (e) S_{10} (60). (.) denotes the angular expansion size.

P_N discretisations. Furthermore, in some of these numerical examples the wavelet bases out performed P_N by some considerable margin. This was particularly highlighted in the final problem for which P_N was unable to provided an approximation that close to the exact solution. The numerical examples also demonstrated the wavelet discretisations to be less susceptible to developing ray effects in their solutions than S_N. However, the demonstrations presented were not able to show any substantial increase in accuracy by using the quadratic functions. In fact, the solutions to the two discretisations were almost identical in all the examples considered. In light of these findings, the linear hexahedral wavelets appear to be the most suitable of the two discretisation methods. This is due to their discretised equations being more efficient to solve as their angular matrices are less dense (of non zero elements) and so require less operations within a linear solver.

Chapter 5

CHEBYSHEV SPECTRAL SPHERICAL WAVELETS FOR ANGULAR DISCRETISATIONS OF THE BOLTZMANN TRANSPORT EQUATION

Contents

This chapter presents a new spherical wavelet method for the angular discretisation of the Boltzmann transport equation. The wavelet basis is based on the principals of second generation wavelets. However, rather than employing the ideas of chapters 3 and 4, and using low ordered linear and quadratic functions over subdivided surfaces, the new wavelet basis uses an alternative method centred on spectral techniques. This spectral wavelet basis utilises Chebyshev collocation Lagrangians and approximate functions over the surface of the hexahedron (which was used in chapter 4 to represent the direction of particle travel). In order to demonstrate the new wavelets capabilities, three 2-D steady state neutron transport problem domains have been solved. The numerical examples aim to demonstrate the discretisation's solutions for accuracy and susceptibility to ray-effects by comparison with the hexahedral wavelets developed in chapter 4. Benchmarked solutions are provided by an established even-parity code.

5.1 Introduction

Over the past 2 decades spectral element methods have seen an increased in their use for discretising and approximating PDEs. Their use have been particularly important to applications involving complex problem domains which, in turn, require approximations of high order accuracy. Some examples of their use include the solving of the Navier-Stokes equations [133, 134, 135, 136, 137, 138, 139, 140], shallow water equations [141], Burgers's equations [142] and advection-diffusion problems [143]. In these examples a

variety of spectral techniques have been used and have shown spectral methods to be particularly powerful for the modelling of fluid flows. In particular, the method use in studying frontal phenomena in oceans has proved rewarding. The methods have also been extensively studied in providing hydrodynamic stability in channel flow problems involving static and variable viscosity. In addition to this, the method can be up to ten times more efficient, in term of memory requirement and CPU time, than a control volume discretisation.

The spectral element method is essentially a high order finite element method that incorporates hp discretisation technologies using bases of polynomial functions. The method inherits the geometric flexibilities of finite elements in which a problem domain is partitioned into a finite number of sub domains (this is a h type discretisation, where h denotes the element scale). The approximation over an element is formed using a polynomial expansion of arbitrary order (this is a p type approximation, where p denotes the polynomial order of the expansion). These properties of the spectral methods make it favourable for approximating demanding problems involving complex geometries; Complex domains are discretised using adaptive element discretisation while the polynomial expansions provide high order accurate solutions.

In this chapter a new spectral, spherical wavelet basis is developed and applied to approximate the angular variable of the Boltzmann transport equation. The basis incorporates spectral method technologies into a second generation wavelet framework. The spectral polynomials are formed through a tensor product of two polynomial expansions which approximate the two directions of the quadrilateral elements that partition the surface of the hexahedron (as with chapter 4 a hexahedron is used to represent the unit sphere and the direction of particle travel). The polynomials employed are the Lagrangian interpolants which correspond to the collocation points positioned over the zeros of the Chebyshev polynomials. The Chebyshev interpolants are then mapped into a hierarchical wavelet basis. This basis is then refined using the lifting scheme in order to increase its stability. The Chebyshev spectral wavelets are designed in order that their function's information concentrates about the vicinity of their collocation point. The nature of the functions are that as the wavelets polynomial order increases, so too does the concentration of their information. Although it will not form part of this chapter's investigation, this property

opens the option for future work of incorporating adaptive angular resolution using these spectral spherical wavelets.

The motivation behind developing these spectral spherical wavelets stems from the fact that the methods can be highly efficiency and accurate for approximation general problems. Analysis shows that the methods are most effective when used to approximate solutions that are smooth. Therefore, the expectation of the technique developed here is that it provides accurate and fast converging solutions - particularly in the case where the angular flux is relatively smooth or isotropic. In addition to this it is also expected that the high order expansion on the sphere will further improve the solutions by means of reducing the severity of ray effects forming in their scalar flux profiles (this is in comparison to the solutions of the lower order wavelets developed in chapter 4). In essence, the improved solutions of the spectral wavelets should more than compensate the increased complexity in their computation. Again this is in comparison to the low order spherical wavelets discussed in chapters 3 and 4.

The following sections are set out as follows. Section 5.2 presents a brief overview of the angular discretised Boltzmann transport equations. In sections 5.3, the spherical spectral wavelets are developed. A demonstration of the new discretisation method is presented in section 5.4 where three steady state neutron transport problems are solved. Section 5.5 finishes this chapter with a conclusion.

5.2 The First Order and Angularly Discretised Transport Equation

This section introduces the angularly discretised, steady state, mono energetic, non multiplying Boltzmann transport equation.

$$A_x \frac{\partial \Psi(r)}{\partial x} + A_y \frac{\partial \Psi(r)}{\partial y} + A_z \frac{\partial \Psi(r)}{\partial z} + \mathcal{H}\Psi(r) = \mathcal{S}(r) \tag{5.1}$$

This is generated using the Bubnov Galerkin weighted residual method, formed from the first order Boltzmann transport equation, where the angular flux $\psi(r, \hat{\Omega})$ is represented by

the set of $\mathcal{M}$ angular basis functions $\mathcal{G}_i$ for $i \in \{1, \ldots, \mathcal{M}\}$.

$$\psi(r, \hat{\Omega}) \approx \sum_{j=1}^{\mathcal{M}} \Psi_j(r) \mathcal{G}_j(\hat{\Omega}) \tag{5.2}$$

The variable Ψ represents the vector of size $\mathcal{M}$ containing the angular moments corresponding to the angular basis functions $\mathcal{G}_i$ for $i \in \{1, \ldots, \mathcal{M}\}$. The matrices A_x, A_y and A_z are $\mathcal{M} \times \mathcal{M}$ angular Jacobians. The $\mathcal{M} \times \mathcal{M}$ matrix $\mathcal{H}$ contains the angular discretised scattering and removal operator and $\mathcal{S}$ is the vector of size $\mathcal{M}$ containing the angular discretised source. These terms are fully defined in chapter 2.

5.3 Chebyshev Collocation Spectral Wavelet Basis on the Hexahedron

This section describes the spectral spherical wavelet basis used to discretise the angular variable of the Boltzmann transport equation. The spherical wavelet basis is essentially a hierarchical extension of the Chebyshev collocation point spectral element method. In the Chebyshev collocation method, a problem domain is partitioned into a finite set of elements, within which a solution (or function) is approximated using a basis of piecewise polynomials. For one dimensional problems, an approximation over an element E_i is obtained by an expansion of (N_i+1) Lagrangian polynomials. The Lagrangians are all of polynomial order N_i and are defined as the set of interpolating functions that correspond to the $(N_i + 1)$ Chebyshev collocation points $\hat{x}_j^i,\ \forall\ j\ \in \{0, 1, \ldots N_i\}$. These are defined on the local element [-1,1] as,

$$\hat{x}_j^i = \cos\left(\frac{\pi j}{N_i}\right), \qquad\qquad j \in \{0, 1, \ldots, N_i\}. \tag{5.3}$$

The approximation of a function $F(x)$, which will be of polynomial order N_i over the local interval, is represented by the expansion,

$$F(x) \approx f(x) = \sum_{j=0}^{N_i} f_j^i\, h_j^i(x), \tag{5.4}$$

where $h_j^i,\ \forall\ j\ \in \{0, 1, \ldots, N_i\}$, denote the Lagrangian functions and $f_j^i,\ \forall\ j\ \in \{0, 1, \ldots, N_i\}$, denote the expansion coefficients.

By definition the Lagrangian functions satisfy,

$$h_j^i(x_k^i) = \delta_{j,k}, \tag{5.5}$$

where $\delta_{j,k}$ denotes the Kronecker-delta function. For the Chebyshev collocation points defined in (5.3) the Lagrangians h_j^i are expressed by,

$$h_j^i(\hat{x}) = \frac{2}{N_i} \sum_{k=0}^{N_i} \frac{1}{c_j c_k} T_k(\hat{x_j})\, T_k(\hat{x}), \tag{5.6}$$

where T_k are the Chebyshev polynomials,

$$T_k(\cos(\theta)) = \cos(k\theta) \tag{5.7}$$

and

$$c_k = \begin{cases} 1, & \text{if} \quad k \notin \{0, N_i\} \\ 2, & \text{if} \quad k \in \{0, N_i\}. \end{cases} \tag{5.8}$$

Figure 5.1 presents the set of collocation points on the local interval for the values $N_i = 1, 2, 4$ and 8. Figure 5.2 presents the corresponding Chebyshev Lagrangian polynomials that interpolate over these positions. In order to expand a known function, the coefficients f_j^i of (5.3) are found by evaluating $F(x)$ over their respective collocation point, $f_j^i = F(x_j^i)$. This ensures that the approximation is exact over all collocation points as illustrated by the expansion of an exponential function shown in figure 5.3. This figure also illustrates the methods spectral accuracy as $N_i = 8$ provides an almost perfect representation.

Chebyshev polynomials possess the important behaviour of not oscillating uncontrollably for high order functions, this is not common among most Lagrangian interpolants. Furthermore, their global maximum are positioned over their collocation points. The functions also concentrate their information about the vicinity of these positions. In fact, the concentration of the function's information becomes more confined to the region surrounding their collocation point as the expansion's polynomial order increases. This is illustrated by the Lagrangian functions corresponding to a 65 collocation point expansion ($N_i = 64$) presented in figure 5.4.

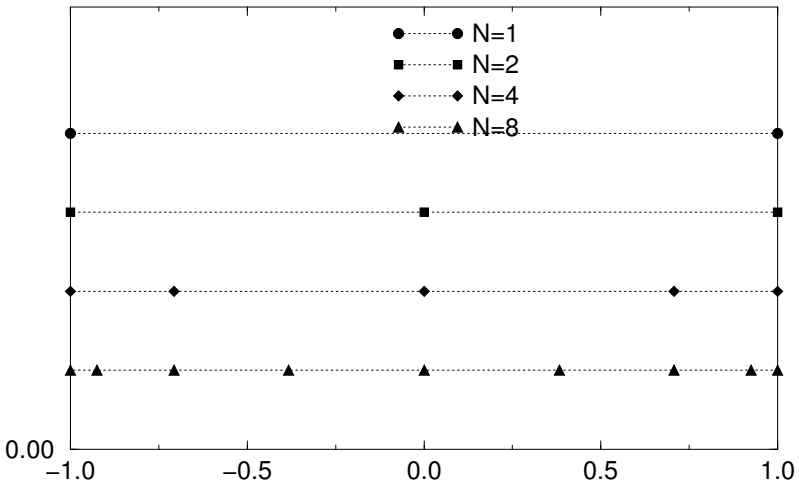

Figure 5.1: The Chebyshev collocation points on the unit interval. The positions displayed are for the values $N = 1, 2, 4$ and 8.

The expansion of a two dimensional problem using Chebyshev collocation polynomials requires the tensor product of two 1-D expansions of the form (5.4). The approximation of a function over the local plane P = [-1,1]×[-1,1] is given by the expansion,

$$f(x, y) \approx \sum_{l=0}^{N_x} \sum_{m=0}^{N_y} f_{l,m}\, h_l(\hat{x})\, h_m(\hat{y}), \tag{5.9}$$

where the functions $h_l(\hat{x})\ h_m(\hat{y}),\ \forall l \in \{1, \ldots, N_x\}, m \in \{1, \ldots, N_y\}$, form a basis for all polynomial functions of up to the order N_x and N_y in the x and y directions, respectively. The terms $f_{l,m}$ denote the expansion coefficients. The basis functions of (5.9) interpolate over the Chebyshev collocation points extended to the local 2-D plane:

$$\{(\hat{x}_l, \hat{y}_m) = (\cos\left(\frac{\pi l}{N_x}\right), \cos\left(\frac{\pi m}{N_y}\right)),\ \ 0 \leq l \leq N_x,\ \ 0 \leq m \leq N_y\}. \tag{5.10}$$

Figure 5.5 shows the set of 2 dimensional collocation points for values $N_x = N_y = 4$ and $N_x = N_y = 8$.

5.3.1 MRA: Chebyshev Lagrangians

A Multiresolution Analysis (MRA) is defined as a nested sequence of subspaces of Lebesgue functional space $L_2(X)$, where X denotes some functional space. In the context of this

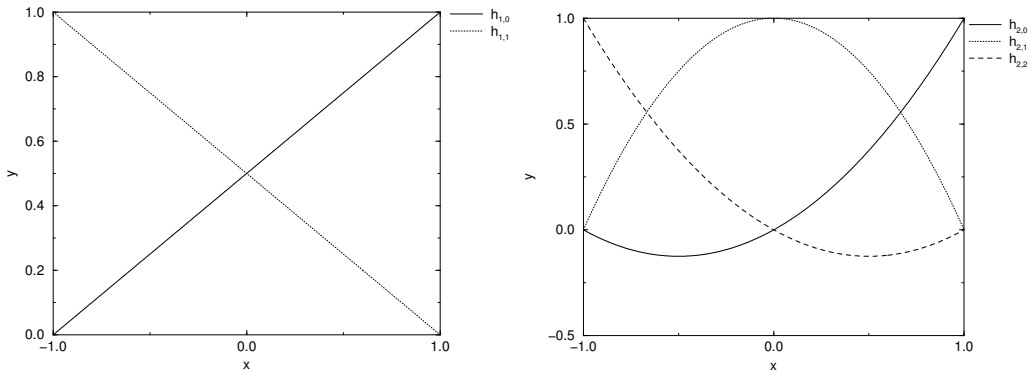

(a) $N = 1$ (b) $N = 2$

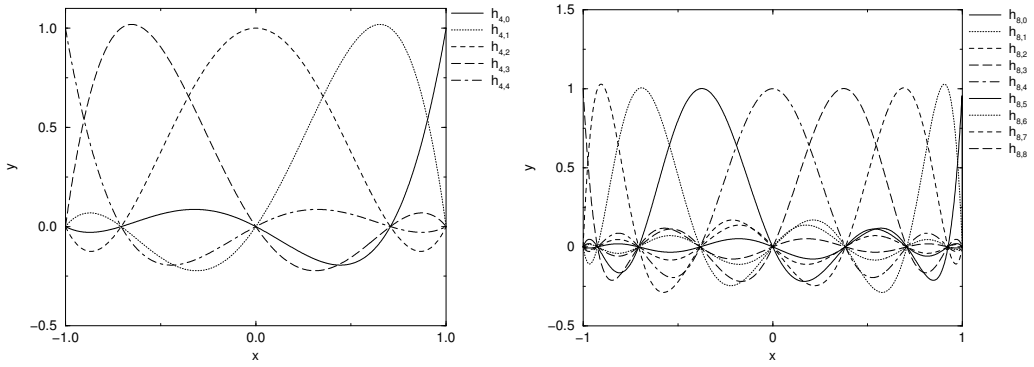

(c) $N = 4$ (d) $N = 8$

Figure 5.2: The Lagrangian functions corresponding to the Chebyshev collocation points for the values: (a) N=1, (b) N=2 , (c) N=4 and (d) N=8.

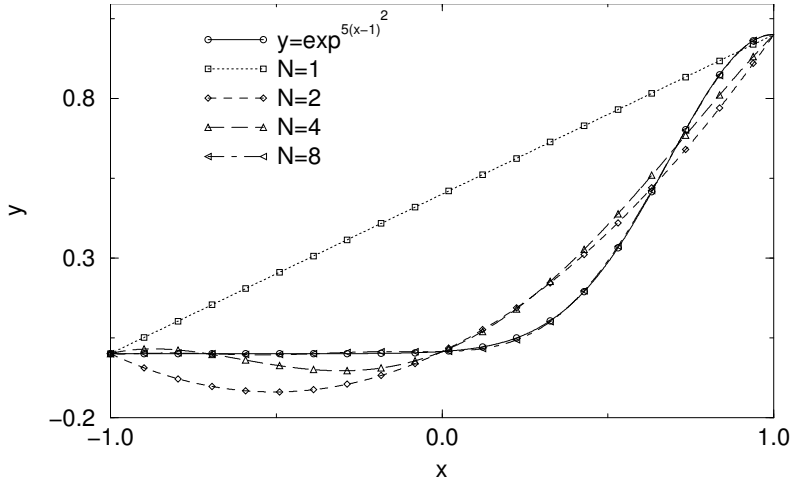

Figure 5.3: An example of the Chebyshev Lagrangian polynomials used in approximating a general function on the unit interval.

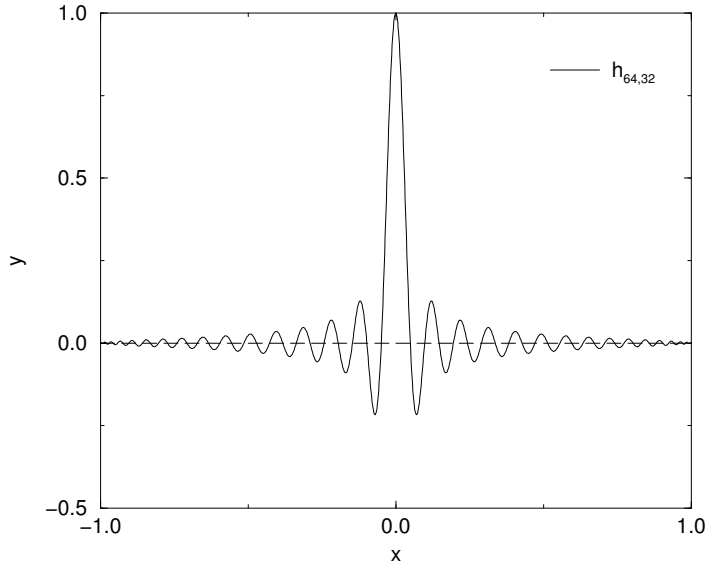

Figure 5.4: A high order Chebyshev Lagrangian polynomial, $h_{64,32}$.

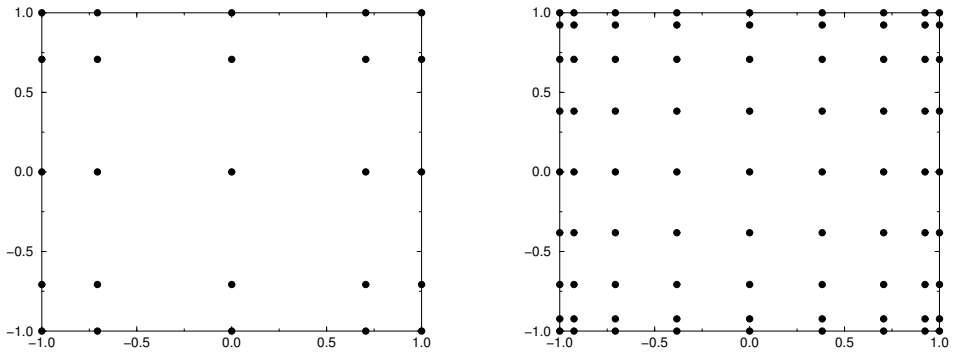

Figure 5.5: A diagram displaying the Chebyshev collocation points on the 2 dimensional unit plane. Left: $N_x = N_y = 4$, Right: $N_x = N_y = 8$, where $N_x + 1$ and $N_y + 1$ denote the number of collocation points placed along the x and y dimensions, respectively.

chapter, X is taken to be the local plane P = [-1,1]×[-1,1] and a function $f \in L_2(\mathsf{P})$ implies that the condition $\int_{\mathsf{P}} f^2 \, dx \, dy < \infty$ holds. Formally, an MRA is defined as a sequence of functional subspaces V_j, $j \in \{0, 1, 2 \ldots \infty\}$, that satisfies the three following conditions,

C1 $V_0 \subset \ldots V_{j-1} \subset V_j \subset V_{j+1} \ldots \subset L_2(\mathsf{P})$

C2 $\bigcup_{j=0}^{\infty} V_j$ is dense in $L_2(\mathsf{P})$

C3 At each level j there exists a set of scaling functions $\varphi_{j,k}, k \in K^{'}(j)$ for some index set $K^{'}$, for which the set $\{\varphi_{j,k}, \;\; k \in K^{'}(j)\}$ forms a Riesz basis of V_j.

The functional spaces V_j are defined here as the set of functions with up to the bi-polynomial order of 2^j over the x and y dimensions of P. For example, V_0, V_1 and V_2 contain the functions with up to bi-linear, bi-quadratic and bi-quartic variation over P, respectively. Conditions C1 and C2 are automatically satisfied by this definition of the spaces V_j. In order to establish a set of basis functions that span the spaces $V_j, \; \forall \; j \; \in \{0, 1, \ldots \infty\}$, the Lagrangians corresponding to the 2-D Chebyshev collocation points on P with $N_x = N_y = 2^j$ are used. This establishes a link between the notation of the MRA and Chebyshev functions. That is, the set $K(j)$ in condition C3 may

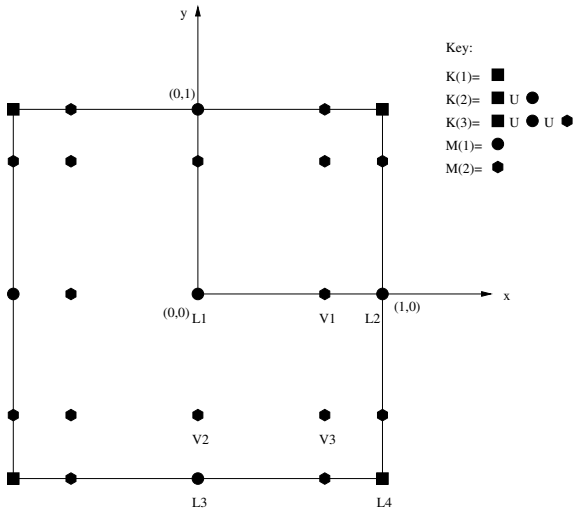

Figure 5.6: A diagram displaying the relationship between the Chebyshev collocation points and the index sets $K(j)$ and $M(j)$ defined in the MRA. The squares denote collocation points of the space V_0, the circles denote collocation points of the space V_1 and the hexagons denote collocation points of the space V_2. Examples of some lazy wavelets and their corresponding lifting functions positions are also shown. The points V1,V2 and V3 correspond to lazy wavelets of the space W_1. Their corresponding lifting functions are (L1,L2), (L1,L3) and (L1,L4), respectively.

index the basis functions of V_j by reference to their unique collocation points,

$$K(j) = \{p = (\hat{x}_l, \hat{y}_m) = (\cos\left(\frac{\pi l}{N_x}\right), \cos\left(\frac{\pi m}{N_y}\right)) \in \mathsf{P},\ \ 0 \leq l, m \leq N_{2^j}\}. \qquad (5.11)$$

The bases for V_j establishes a nested structure within the collocation points. That is, the positions indexed by $K(j)$ are also indexed by the set $K(j+1)$, $\forall\ j \in \{0, 1, \ldots, \infty\}$. This was illustrated in figure 5.1 showing the first 4 sets of collocation points for the 1-D example. The nested points allow a second index set $M(j)$ to denote the new collocation points formed in the V_{j+1} basis, i.e. $M(j) = K(j+1) \setminus K(j)$ where '\' is the set removal operator. An illustration of these index sets $K(j)$ and $M(j)$, for $j \in \{0, 1, 2\}$, is presented in figure 5.6.

A function F defined on P can be approximated by its projection onto a subspace V_j of the MRA. The projected function f_j is then represented by an expansion of the basis

functions of V_j,

$$F(x,y) \;\approx\; f_j \;= \sum_{l=0}^{2^j} \sum_{m=0}^{2^j} f_{l,m}\, h_l(\hat{x})\, h_m(\hat{y}), \tag{5.12}$$

where the functions $h_l h_m$ are defined in (5.6) and (5.9). Using a suitable re-ordering, the expansion is re-written into a single summation form that conforms with the notation of the MRA,

$$F(x,y) \;\approx\; f_j \;= \sum_{k=1}^{(2^j+1)\times(2^j+1)} \lambda_{j,k}\, \varphi_{j,k}. \tag{5.13}$$

The functions $h_l(\hat{x})\; h_m(\hat{y})$ and expansion coefficients $f_{l,m}$ are now represented by $\varphi_{j,k}$ and $\lambda_{j,k}$, respectively. The general properties of an MRA and their associated scaling functions have not been included in this chapter as an overview of the main details have been covered in chapters 3 and 4. A full description can also be found in [94] and [144].

5.3.2 A Spectral Wavelet Basis on the Local Element P

Wavelets [132, 94] are functions that have a general purpose to form bases for the space W_j which complement the functional spaces V_j in V_{j+1}, $\forall j \in \{0, 1, \ldots, \infty\}$. That is, the space V_{j+1} can be decomposed as $V_{j+1} = V_j \oplus W_j$, $\forall j \in \{0, 1 \ldots, \infty\}$. By recursively applying this decomposition of the functional spaces of an MRA, a space V_j is decomposed using the spaces V_0 and W_l, for $l \in \{0, 1, \ldots j-1\}$.

$$V_j = V_0 \oplus_{l=0}^{j-1} W_l \tag{5.14}$$

By letting $j \to \infty$ the decomposition (5.14) generates a basis for the space $L_2(\mathsf{P})$. Also, following from this decomposition, an alternative expansion to the approximation in equation (5.13) is generated,

$$f_j = \sum_{k \in K(0)} \lambda_{0,k}\, \varphi_{0,k} + \sum_{l=0}^{j-1} \sum_{k \in M(l)} \gamma_{l,k}\, \psi_{l,k}. \tag{5.15}$$

The functions $\varphi_{0,k}$ are the scaling functions that form the basis for the lowest order MRA space V_0. The terms $\lambda_{0,k}$ are the scaling function's expansion coefficients. The functions $\psi_{l,k}$ are the wavelets that form the bases of the complementing spaces W_l, for $l \in \{0, 1 \ldots, j-1\}$. These wavelets are formally defined as,

$$\{\psi_{l,k}, k \in M'(l)\}, \tag{5.16}$$

where $M^{'}(l)$ denotes a set indexing the wavelets. The wavelet expansion coefficients of (5.15) are denoted by $\gamma_{l,k}$.

The bases for the complementing spaces W_l are built using the techniques used in the construction of second generation wavelets. This process requires two steps: Step 1 - Construct a simple 'lazy wavelet' basis for W_l that complements V_l in V_{l+1}, $\forall\ l\ \in \{0, 1, 2, \ldots, \infty\}$. Step 2 - modify the wavelets of W_l, $\forall\ l \in \{0, 1, 2, \ldots, \infty\}$, into more stable functions using the lifting method.

The lazy wavelet bases are generated using the nested property of the collocation points. A simple wavelet basis that spans the complementing space W_l is defined by

$$\psi_{l,m} = \varphi_{l+1,m}, \qquad \forall\, m\ \in\ M(l), \quad \forall\, l\ \in\ \{0, 1, \ldots, \infty\}. \tag{5.17}$$

The fact that (5.17) generates bases which span W_l, $\forall\ l\ \in\ \{0, 1, 2, \ldots, \infty\}$, is easily proved.

Theorem:
The spaces W_j spanned by the lazy wavelets $\psi_{j,m}$ defined in (5.17) complement the spaces V_j in V_{j+1} for all $j \in \{0, 1, 2, \ldots \infty\}$.

Proof:
By showing that all basis functions of the space V_{j+1} can be generated by the functions $V_j \oplus W_j$, and showing that the converse also holds, the theorem will be proved since the basis functions for each space are contained within the span of the other space, i.e. $\varphi_{j+1,m} \in V_j \oplus W_j$ and $\varphi_{j,m}, \psi_{l,k} \in V_{j+1}$, therefore span$\{\varphi_{j+1,m}\}$ = span$\{\varphi_{j,m}, \psi_{l,k}\}$.

First consider any scaling function $\varphi_{j+1,k} \in V_{j+1}$, for $k \in K(j+1)$, that is associated with the collocation point $p_k \in \mathsf{P}$, then, either $k \in K(j)$ or $k \in M(j)$. For $k \in M(j)$, the result follows from (5.17) where by definition, there exists a lazy wavelet in W_j identical to $\varphi_{j+1,k}$.

For $k \in K(j)$, a function $\varphi^{'}$ can be generated by a combination of the scaling functions and wavelets of V_j and W_j, respectively:

$$\varphi^{'} = \varphi_{j,k} - \sum_{m \in M(j)} \varphi_{j,k}(p_m)\psi_{j,m}. \tag{5.18}$$

The first contributing term in (5.18) corresponds to the scaling function in V_j, associated with the collocation point p_k. The summation then removes a proportion of the lazy wavelets from the function $\varphi_{j,k}$. The contribution removed by each wavelet $\psi_{j,m}$, $m \in M(j)$, is the value of the scaling function $\varphi_{j,k}$ evaluated at the collocation point p_m, $m \in M(j)$. The function $\varphi^{'}$ is therefore of bi-polynomial order 2^{j+1} over the x and y coordinates. In addition to this, the construction of $\varphi^{'}$ ensures the function has the value of 1 over the collocation point p_k (indexed by $k \in K(j+1)$) and zero over all other collocation points in $K(j+1)$. These are properties are also shared with the scaling function $\varphi_{j+1,k}$. Therefore, since there exists only one function of bi-polynomial order of 2^{j+1} that possesses these values over the $(2^{j+1}+1)^2$ collocation points, the two functions $\varphi_{j+1,k}$ and $\varphi^{'}$ must be identical.

The converse is simple to show. For any wavelet function $\psi_{j,k}$, for $k \in M(j)$, by definition there exists a function $\varphi_{j+1,l}$, for $l \in K(j+1)$, such that $\psi_{j,k} = \varphi_{j+1,l}$. Also, for any scaling function $\varphi_{j,l}$, for $l \in K(j)$, then, by definition of the MRA, $\varphi_{j,l} \in V_j \subset V_{j+1}$ which implies that $\varphi_{j,l} \in V_{j+1}$. This completes the proof.

The wavelets defined in (5.17) are modified in order to generate a more stable Riesz basis. For this, the modification process lifting is used. In the lifting method, contributions from the scaling functions of the space V_j are removed from the lazy wavelets $\psi_{j,k}$, for $k \in M(j)$,

$$\psi_{j,m} = \varphi_{j+1,m} - \sum_{k \in K(j)} S_{j,m,k}\, \varphi_{j,k} \qquad m \in M(j) \tag{5.19}$$

The removed contributions from the lazy wavelet, denoted by $S_{j,m,k}$, are not restricted as the condition $V_j \oplus W_j = V_{j+1}$ holds for arbitrary values. The lifting contribution is therefore selected in order that the wavelets possess two important properties: zero integrals and locality. Wavelets possessing zero integrals have been demonstrated to provide more stable Riesz bases [119]. Locality is required in order to apply local resolution on the sphere for future adaptive angular discretisation methods in approximating the BTE. Although the spectral function's supports cover the whole space P, it was shown in a previous section that each function's information is concentrated about its collocation point. Therefore, in order to keep the lifted wavelets information local to their collocation point, only the two closest neighbouring scaling function to this position are used in its lifting.

Figure 5.6 shows the collocation positions used in the lifting process. In order to satisfy these two conditions, the lifting weights are defined as,

$$S_{j,k,m} = \begin{cases} \frac{\int \varphi_{j+1,m}}{2 \int \varphi_{j,k}} & \text{if } m \in M(j) \text{ is adjacent to } k \in K(j) \\ 0 & \text{otherwise.} \end{cases}$$

Figure 5.7 presents a lazy wavelet of the space V_2, its two lifting functions and the lifted wavelet.

5.3.3 Representation of Spherical Functions

At present, the wavelets approximate functions on quadrilateral elements mapped to the local plane P. In order for the basis to represent spherical functions, 24 quadrilateral elements are positioned to construct a hexahedral surface as shown in figure 5.8. Spherical functions are then approximated on the surface of the hexahedron using the mapping between the two domains shown in figure 5.9. The spherical wavelets can approximate with continuity or discontinuity over the elements' boundaries. The former method being achieved by combining wavelets on adjacent elements that share the same collocation positions. Note, however, discontinuity along the planes $x = 0$, $y = 0$ and $z = 0$ is required in order to avoid ray effects dominating solutions of 1-D and 2-D neutron transport problems [132].

5.4 Numerical Examples

This section demonstrates the spectral hexahedral wavelets capabilities by solving three 2-D steady state neutron transport problems. The problems have been selected in order to demonstrate the angular discretisation's accuracy and their susceptibility to forming ray-effects in their solutions. The numerical examples are also designed to establish the benefits of using spectral methods over the low order linear wavelet functions.

For the following calculations the spatial variables are discretised using an SUPG control volume finite element method [88]. Quadrilateral bi-linear elements are employed to dis-

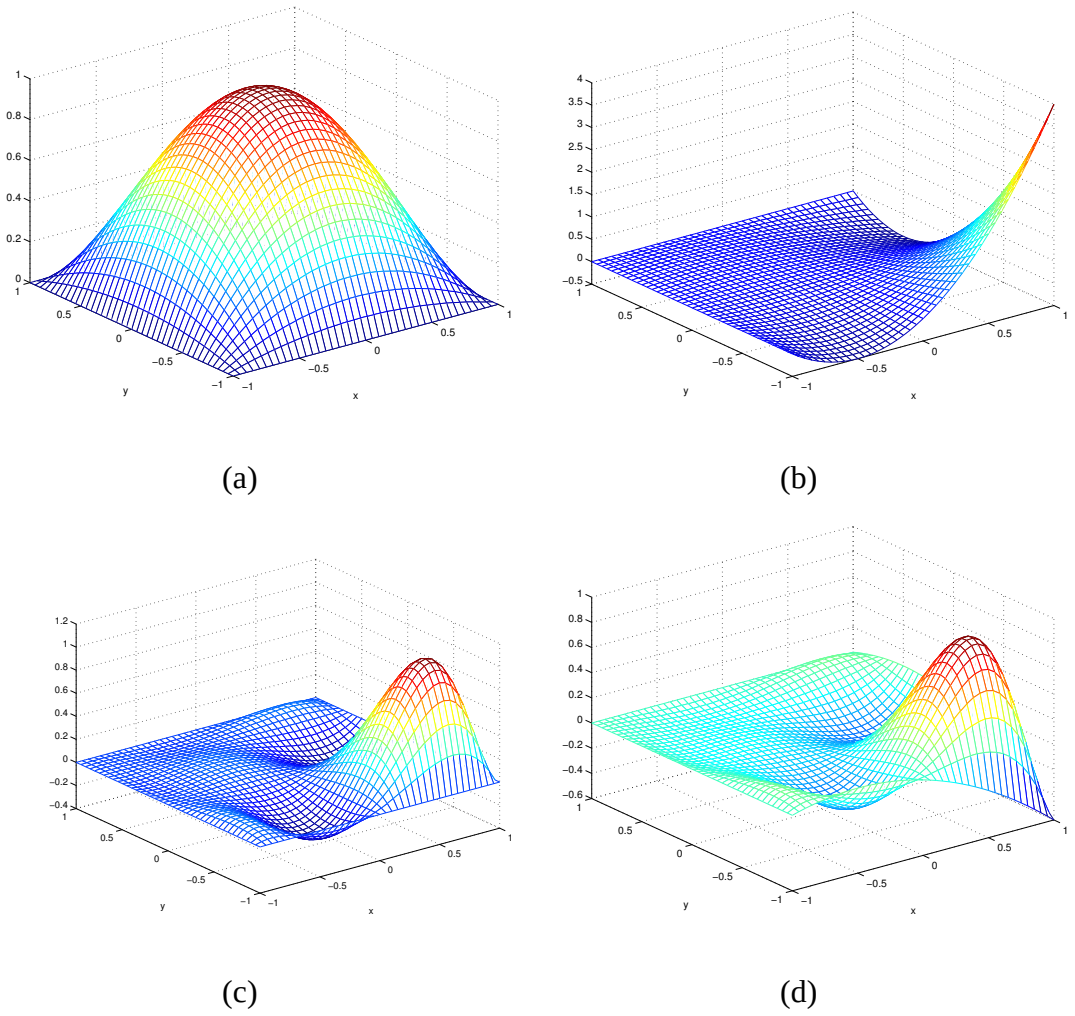

Figure 5.7: Examples of some Chebyshev scaling functions and wavelets on the unit plane. (a) The scaling function of the space V_1 corresponding to the mid-plane collocation point (point L1 in figure 5.6). (b): The scaling function of the space V_1 corresponding to the corner node collocation point (point L4 in figure 5.6). (c): The scaling function of the space V_2 corresponding to the collocation point V3 in figure 5.6 (this is also a lazy wavelet of the space W_1). (d): The lifted wavelet of the lazy wavelet in figure (c) (the lifting functions are the two functions shown figures (a) and (b)).

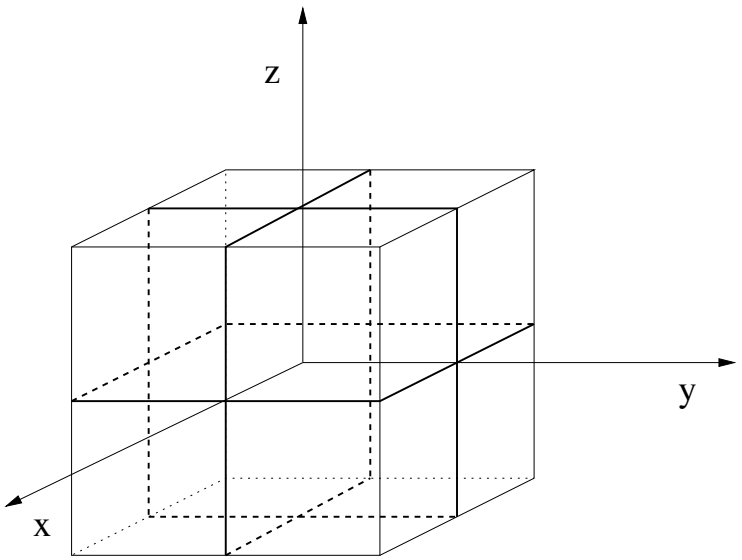

Figure 5.8: The spectral wavelets approximate over the 24 local elements that are used to construct the surface of the hexahedron. The approximation will allow continuity or discontinuity across the element boundaries. However, discontinuity is required across the planes $x = 0$, $y = 0$ and $z = 0$ (outlined in bold) in order to avoid ray effect type inaccuracies.

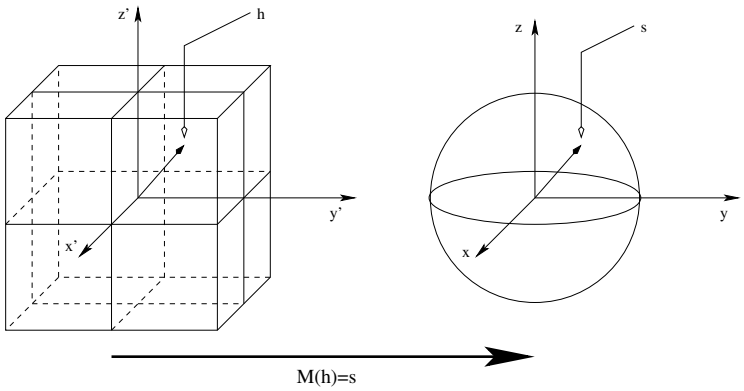

Figure 5.9: A diagram displaying the mapping from the hexahedron to the unit sphere to enable the representation of spherical functions. This is simply the normalising mapping $s = (x, y, z) = M(h) = (x^{'}, y^{'}, z^{'})/(x^2 + y^2 + z^2)^{\frac{1}{2}}$, where $s \in S^2$ and $h \in \mathsf{H}$

Angular Expansion sizes		
N (Angular expansion)	HLW_N	HCW_N
0	28	28
1	76	76
2	244	244
3	868	868

Table 5.1: This table lists the angular expansion sizes of the hexahedral linear and spectral wavelet discretisations.

cretise and approximate the problem's spatial domain. The Chebyshev wavelet solutions relating to the space V_N (over all elements of the hexahedron H) is denoted by HCW_N. The expansion sizes of the wavelet discretisations are listed in table 5.1. Benchmarked solutions are provided by the established even-parity code EVENT.

5.4.1 Problem 1: A Box Problem Containing a Pure Absorbing Medium

This problem is designed to demonstrate the spectral wavelets capabilities on a simple pure absorbing problem. The aims are compare the accuracy of the Chebyshev wavelets and to investigate their capabilities in comparison to the linear Hexahedral wavelets described in chapter 4.

The problem domain is presented in figure 5.10 and its corresponding material and source information is listed in table 5.2. The domain is a 3cm by 3cm square containing a material with a constant pure absorbing cross-section of $\sigma_a = 0.5\mathrm{cm}^{-1}$. The central 0.5cm by 0.5cm square contains an isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and vacuum boundaries are prescribed along all edges of the problem domain. The following calculations were performed using a structured quadrilateral grid of 225 elements and 256 nodes, see figure 5.11. The benchmarked solution is provided by EVENT using a P_{23} angular approximation and a linear, finite element discretisation of the spatial domain using a fine spatial mesh. This solutions is presented in figure 5.12.

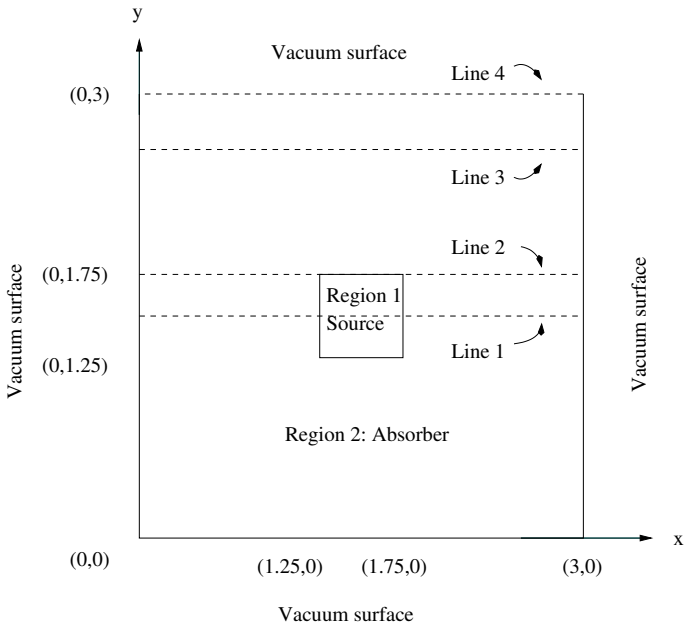

Figure 5.10: Problem 1: The diagram of the box problem.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.0
2	0.0	0.5	0.0

Table 5.2: Source and material properties of Problem 1

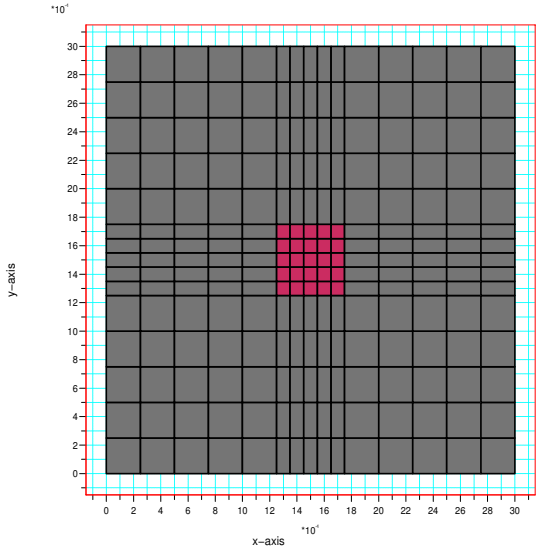

Figure 5.11: Problem 1: The spatial mesh of quadrilateral finite elements used for the spatial discretisation in all wavelets calculations.

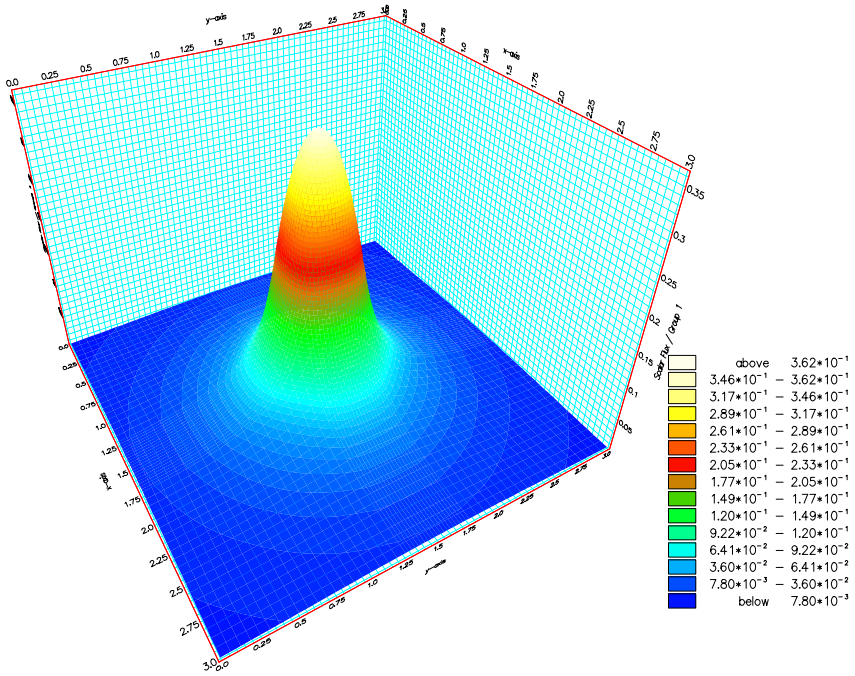

Figure 5.12: Problem 1: The benchmarked solution provided by EVENT P_{23}.

Figure 5.13 presents the scalar flux solutions of HLW_2 and HCW_2. Both solutions display similar likenesses, they peak in flux over the centre of the domain at approximately 0.36 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and attenuate sharply over the outer regions of the source. The flux then enters a gradual decline towards the boundary of the problem domain. These profiles are in close agreement with the exact solution provided by EVENT, figure 5.12. Note that the smoother profile of the exact solution arises through the fine spatial mesh used. Therefore the courser profiles displayed in the wavelet solutions should not be perceived as a limitation of the angular approximation, but as a consequence of the courser spatial mesh.

Figure 5.14 presents the scalar flux contours of the HLW_0, HLW_1, HLW_2, HCW_1, HCW_2 and exact solutions (note that HLW_0 and HCW_0 are identical discretisations). The results show the high order linear and spectral wavelets solutions to be near angular convergence and that they have mitigated the large oscillations (ray effects) which formed in the HLW_0 (and HLW_0) solution. The results also show few visible ray-effects have formed in the solutions of the HLW_1 and HCW_1 discretisations. However, these approximations have formed some small un-physical fluctuations over the horizontal and vertical direction from the source (where the solutions exhibit a slightly increased scalar flux). For equal expansion sizes, the linear and spectral wavelets solutions appear to be in very close agreement.

Figure 5.15 presents the 4 higher discretisation's scalar flux solutions through the lines 1-4 as shown in figure 5.10. The first graph plots the flux through the centre of the domain. The graph shows all wavelet solutions to be in close agreement with the benchmarked solution. This is due to the central region producing fairly isotropic fluxes which, in turn, is easily approximated by the low order wavelet discretisations. This therefore explains the reason for seeing only minor differences in the scalar flux when the wavelet approximations are increased.

A similar result is seen in the graph plotting the flux along line 2. Again, the reasonably high isotropic distribution of the flux along these positions has led to all wavelet solutions being in close agreement. The graph also shows the wavelet solutions are not in close agreement with the Benchmarked solution. However, these discrepancies are a product of

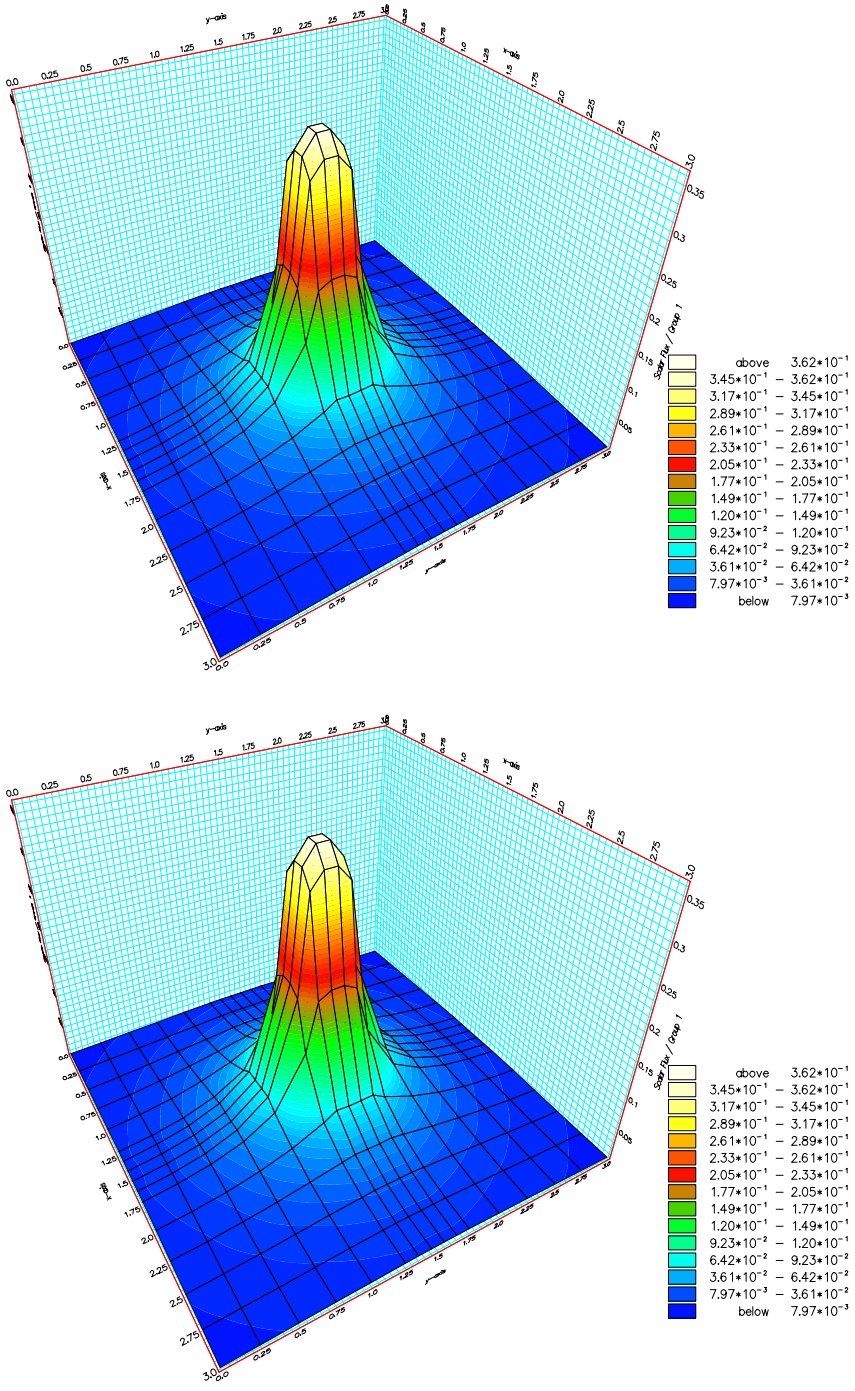

Figure 5.13: Problem 1: Scalar flux profiles of the HLW_2 (top) and HCW_2 (bottom) solutions.

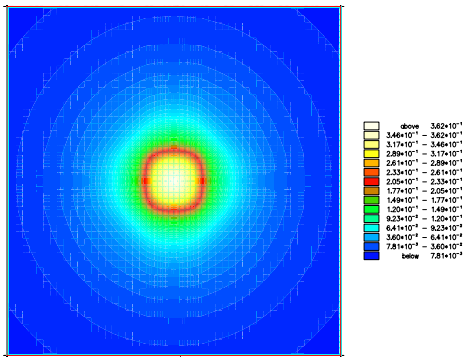

(a) Benchmark

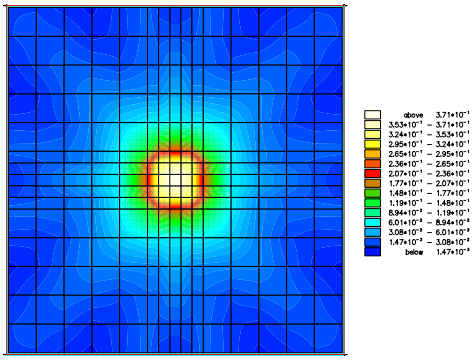

(b) HLW_0/HCW_0

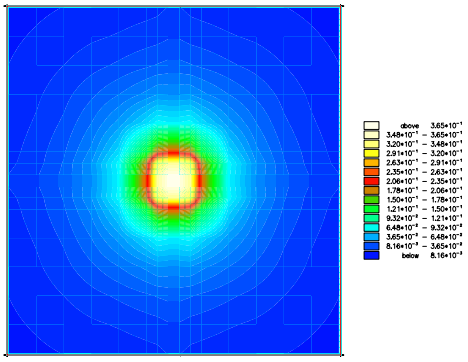

(c) HLW_1

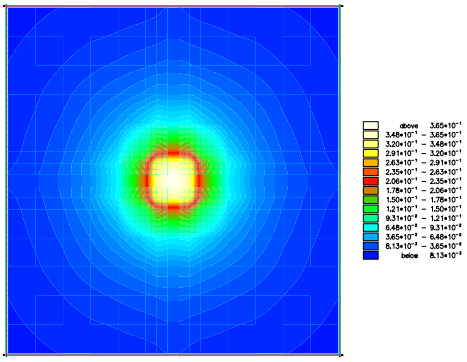

(d) HCW_1

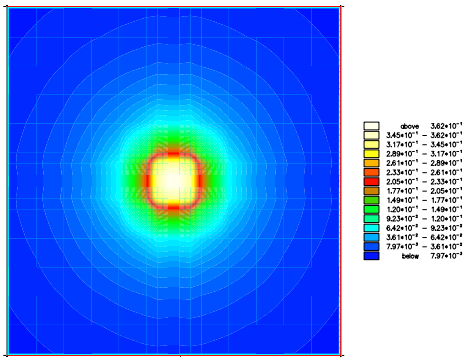

(e) HLW_2

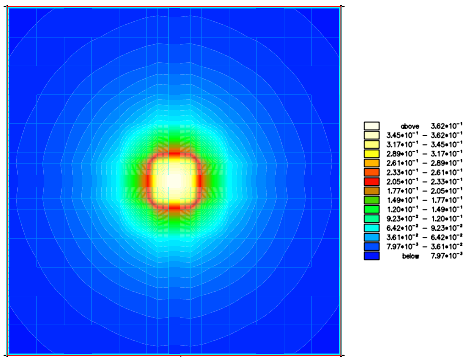

(f) HCW_2

Figure 5.14: Problem 1: The scalar flux contour solutions. (a) Benchmarked solution, (b) HLW_0/HCW_0 (28), (c) HLW_1 (76), (d) HCW_1 (76), (e) HLW_2 (244), (f) HCW_2 (244). The term (.) denotes the number of angular basis functions used in the angular approximation.

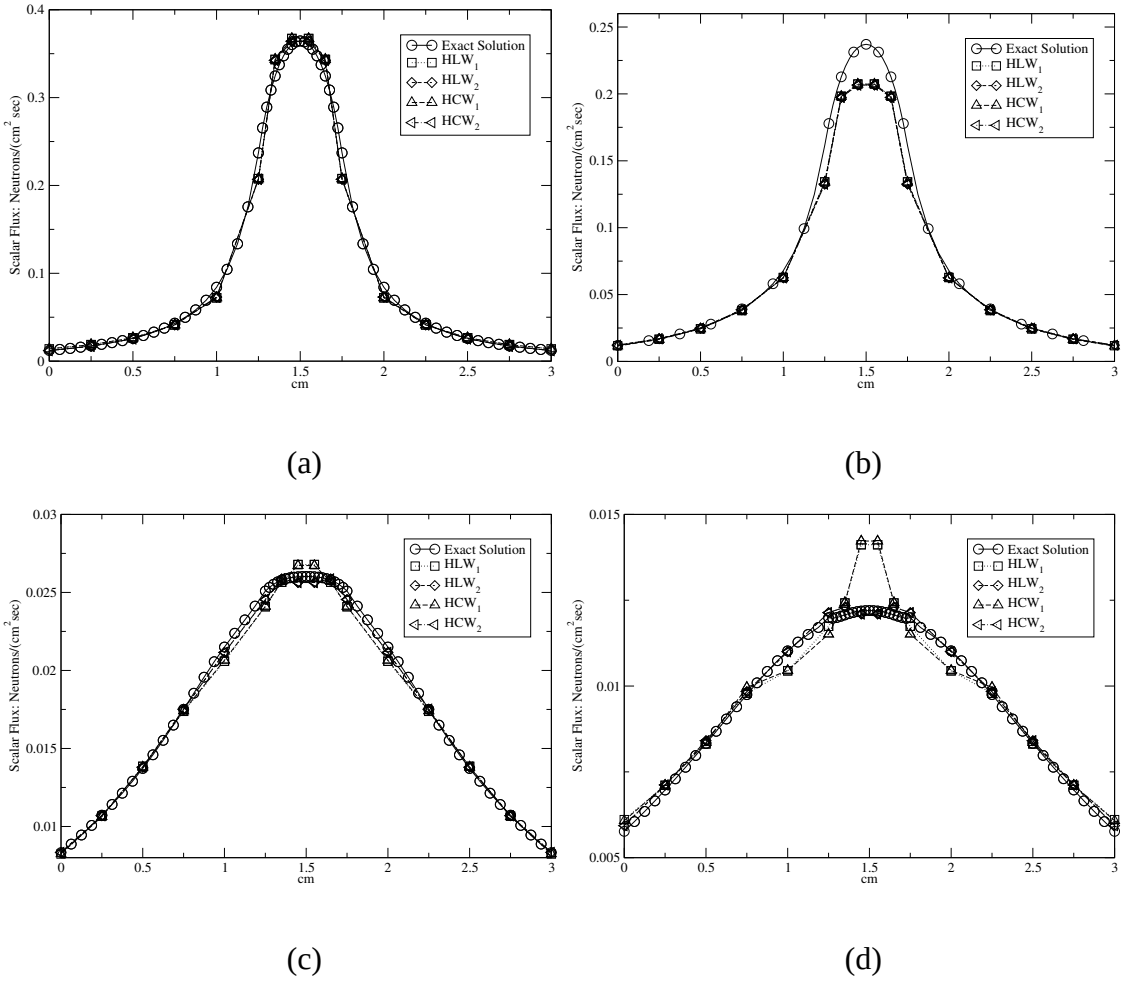

Figure 5.15: Problem 1: Scalar flux profiles of the benchmarked, HLW_1, HLW_2, HCW_1 and HCW_2 solutions. The line cut positions are as follows. (a) (0,1.5),(3,1.5), (b) (0,1.75),(3,1.75), (c) (0,2.25)(3,2.25), (d) (0,3)(3,3).

the course spatial approximation as it can be seen from the profiles presented in figure 5.13 the scalar flux experiences high gradients of change. Therefore, the course spatial mesh will show its highest susceptibility to errors over these positions. The close agreement of the wavelet solutions also suggests that angular convergence has been reached using this spatial discretisation.

The graph presenting the scalar flux across line 3 reveals a number of variations in the wavelet's solutions. The result shows the peak in scalar flux is more accurately represented using the high order HLW_2 and HCW_2 wavelet discretisations. The lower order HLW_1 and HCW_1 discretisations generate a maximum in the scalar flux that reaches 0.0268 and 0.0267 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, respectively. This is in comparison to the benchmarked solution of 0.0260 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The higher ordered linear and spectral wavelets are shown to be almost identical. With both peaking at 0.0256 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ they reduce the error of the lower ordered wavelets by approximately 50%. The higher order approximations are also shown to have improved the approximation over the intervals (0.75,1.35) and (1.65,2.25). In these regions the lower ordered solutions display ray-effects that cause their solution to oscillate about the exact solution.

The greatest improvements of using the high ordered wavelets are seen in the result plotting the leakage through the top boundary. The peak at the centre of the boundary is over approximated by a considerable margin using the low ordered wavelet discretisations. In comparison with the benchmarked solution, which measured 0.0122 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, the low ordered wavelets peaked above this value at 0.0142 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Increasing the expansions to the higher ordered wavelets give significant improvements. Both discretisations resolve the peak closer to the benchmarked solution at 0.0121 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Ray-effect oscillations are also clearly visible in the scalar flux solutions along the problem's boundary. In both low ordered calculations, ray-effects have caused large oscillations about the exact solution. These oscillations account for the high peaks in the graph's centre and the low troughs in the regions either side of the maximum. The high ordered wavelets have suppressed most of these ray-effect distortions, however, small oscillations about the exact solution are still visible over the interval (1.25,1.75).

In order to establish the spectral wavelet's ability to increase the solutions accuracy, the

scalar flux error norms are evaluated for each wavelet calculation. The error norms of the scalar flux solutions are defined as,

$$\text{error}(\phi_{approx}) = (\int_V (\phi_{approx} - \phi_{exact})^2 dV)^{\frac{1}{2}} \tag{5.20}$$

where ϕ_{approx} and ϕ_{exact} denote the approximate and exact scalar flux, respectively. (Note that in order to make the spatial discretisation consistent, RADIANT is used to calculate the exact scalar flux. This was obtained using the highest order angular discretisation possible to ensured the angular approximation is the closest we can get to being converged.)

Figure 5.16 presents the graphs of the error norms against the square root of the angular expansion size. The graph shows the Chebyshev wavelets to have the greatest order of accuracy - although the difference between the methods is not as large as would be expected. One possible reason for this may be due to the problem being simple enough to approximate accurately using just the linear wavelets. Therefore, increasing the spectral representation only leads to a small increases in accuracy. It is also important to note that although the exact solution is very accurate, it is still an approximation. Therefore one must bare in mind that the small scalar flux errors seen in this graph may be influenced more by the inaccuracies of the exact solution.

5.4.2 Problem 2: A Box Problem Containing an Absorbing and Scattering Medium

This problem is designed to demonstrate the wavelets accuracy for resolving domains involving materials with isotropic scattering cross-sections. For this investigation the wavelet solutions are compared to the hexahedral wavelets detailed in chapter 4.

The problem domain is presented in figure 5.17 and the corresponding source and material information is listed in table 5.3. The problem domain is a 4cm by 4cm square. The central 2cm by 2cm square contains an isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The problem contains a material with an absorbing cross-section $\sigma_a = 0.6\ \mathrm{cm}^{-1}$ and a scattering cross-section $\sigma_s = 0.4\ \mathrm{cm}^{-1}$. Vacuum boundary conditions are prescribed across all edges of the problem. The following calculations were performed using a struc-

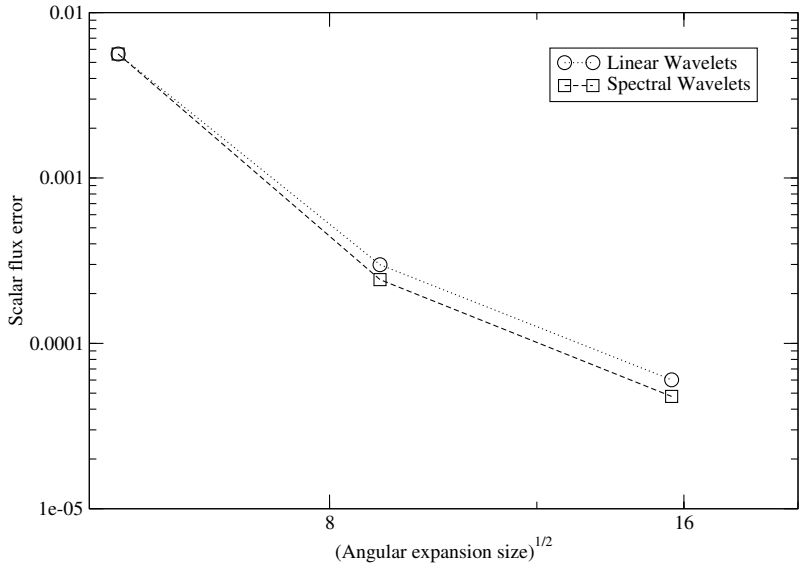

Figure 5.16: This graph presents the scalar flux error norms of the linear and spectral wavelet discretisations of problem 1 - the box source problem

Region	Source $(cm^{-2}s^{-1})$	$\sigma_a\ (cm^{-1})$	$\sigma_s\ (cm^{-1})$
1	1.0	0.6	0.4
2	0.0	0.6	0.4

Table 5.3: Source and material properties of Problem 2: The Box Source Problem 2.

tured, quadrilateral finite element grid containing 400 elements, of size 0.2×0.2cm, and 441 nodes. The mesh is presented in figure 5.18. The benchmarked solution was provided by EVENT using a P_{23} angular approximation and a linear, finite element (fine mesh) discretisation of the spatial variables.

Figure 5.19 presents the benchmarked, HLW_1 and HCW_1 solutions. All three solutions appear to be in close agreement. Their scalar flux profiles peak over the domain's centre at approximately 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and attenuate quickly towards the problem's boundaries.

In order to demonstrate the wavelet discretisation's accuracy and susceptibility to forming ray-effects, the remaining analysis of this problem concentrates on the flux leakage

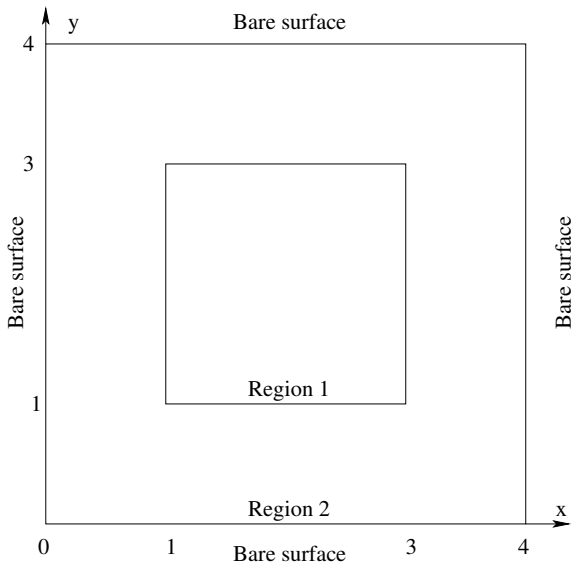

Figure 5.17: Problem2: The diagram of the box problem.

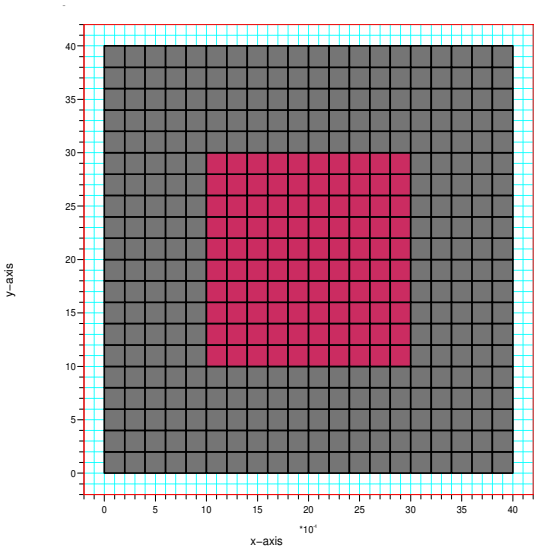

Figure 5.18: Problem 2: The spatial mesh of quadrilateral elements used for the spatial discretisation for all wavelets calculations.

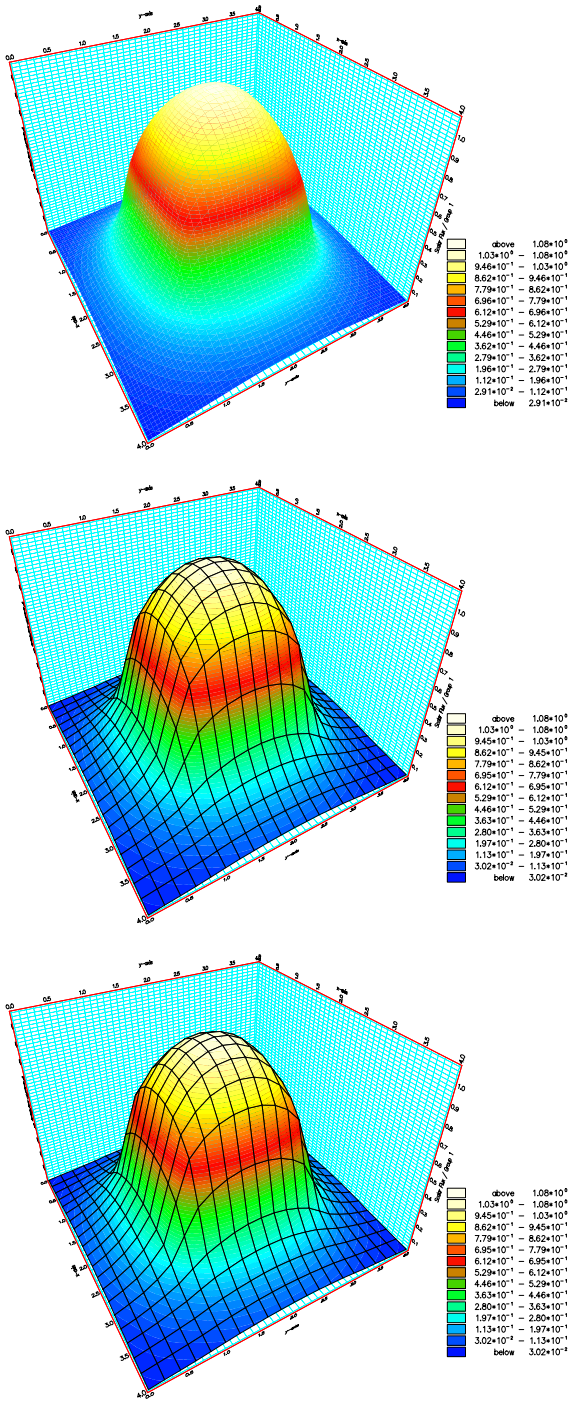

Figure 5.19: Problem 2: Scalar flux profiles of the exact (top), HLW_1 (middle) and HCW_1 (bottom) solutions.

through the top boundary of the problem - where most errors are likely to occur. Figure 5.20 presents two graphs that plot the wavelet's scalar flux solutions from position (0,4) to (2,4) across the top boundary.

The top graph presents the HLW_0 (or HCW_0), HLW_1 and HCW_1 scalar flux solutions. The HLW_0 (or HCW_0) discretisation forms a number of ray effects that cause its solution to oscillate about the benchmarked solution. The largest of these oscillations occur over the positions $x = 0.4$ and $x = 1.6$. These fluctuation are significantly reduced using HLW_1 and HCW_1 (for which both discretisation's solutions follow similar flux profiles). The solutions are in reasonable agreement with the benchmarked solution over the region $x \in (0, 0.8)$. However, small oscillations still persist to distort the solution across the remaining regions.

The bottom graph presents the HLW_2 and HCW_2 scalar flux solutions. These results show no sign of oscillations about the exact solution. The two wavelet solutions are again in very close agreement and display a higher than expected peak in the scalar flux at $x = 2$ (this was found to also have occurred in the HLW_1 and HCW_1 approximations). The wavelet solutions are therefore showing signs of angular convergence. This demonstrates the over approximated flux at $x = 2$ to be a product of the low resolution spatial discretisation.

5.4.3 Problem 3: A void duct surrounded by a pure absorber

This problem is designed to demonstrate the spectral wavelet's ability to resolve neutron streaming through voids. For this investigation the Chebyshev wavelets are compared with the solutions of the Hexahedral wavelets of chapter 4.

The problem domain is presented in figure 5.21 and the corresponding source and material information is listed in table 5.4. The problem has a rectangular domain of size 36cm $\times$ 28cm. The central 6cm $\times$ 6cm square region contains an isotropic source of strength 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and a purely absorbing material with a cross-section $\sigma_t = 0.5\mathrm{cm}^{-1}$. Two ducts containing voids of width of 6cm extend from the source to the problem's upper

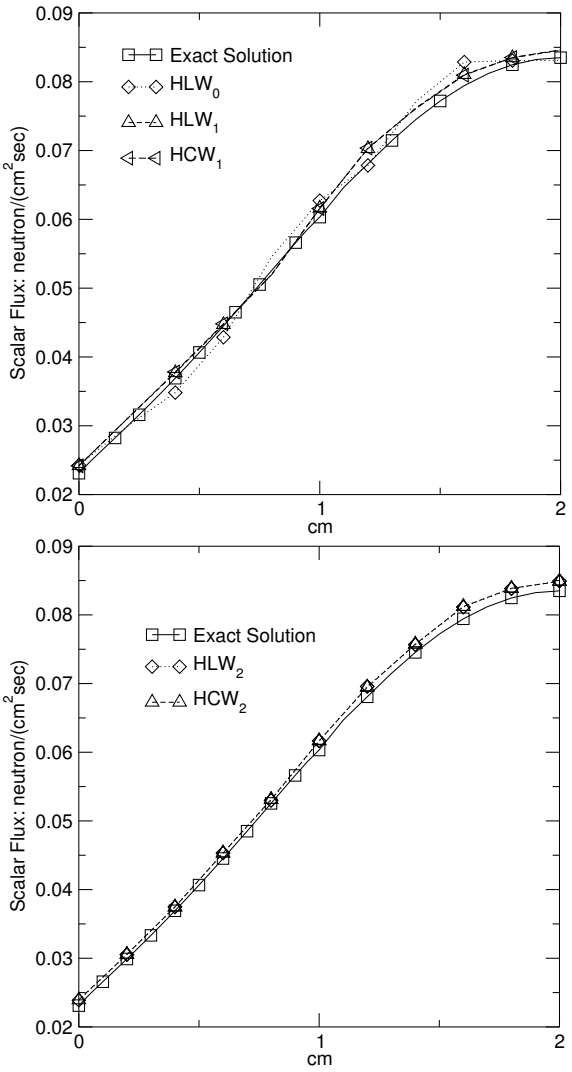

Figure 5.20: Problem 2: The scalar flux leakage through the top boundary of the problem domain, that is, through the points (0,4),(2,4). Top graph: The HLW_0/HCW_0, HLW_1 and HCW_1 solutions. Bottom graph: The HLW_2 and HCW_2 solutions.

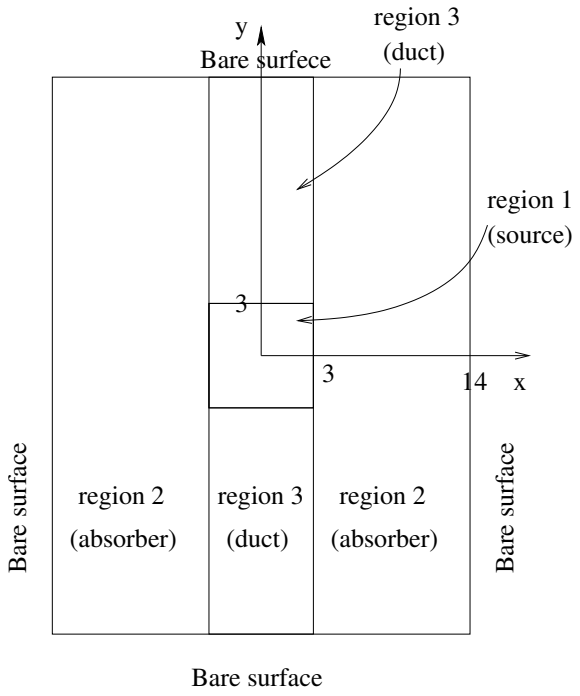

Figure 5.21: Problem 3: This diagram displays the domain of the straight duct problem.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.0
2	0.0	0.5	0.0
3	0.0	0.0	0.0

Table 5.4: Problem 3: Source and material properties of the Straight Duct problem

and lower boundaries. The material surrounding the ducts has a absorbing cross-section $\sigma_t = 0.5\mathrm{cm}^{-1}$. Vacuum boundaries are prescribed along all edges of the problem. A finite element mesh consisting of 3,600 elements was employed to discretise the problem's spatial domain.

Figure 5.22 presents the scalar flux solution using RADIANT set to the highest order angular approximation available. The solution is therefore as near to angular convergence that RADIANT is capable of achieving on the 3,600 finite element spatial mesh. This is therefore taken as the benchmark solution. The scalar flux solution peaks in the problem's

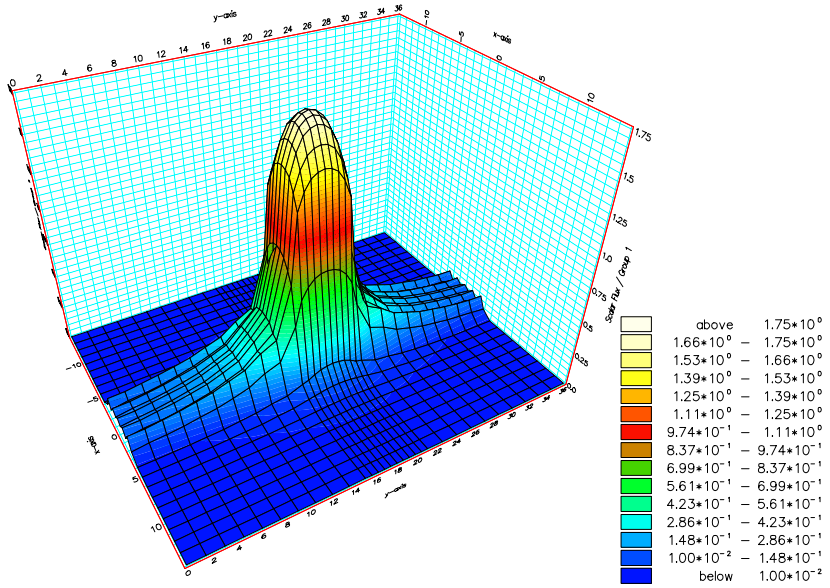

Figure 5.22: Problem 3: The angularly converged scalar flux solution of the 3,600 finite element spatial discretisation mesh.

centre at approximately 1.7 neutrons $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. In the absorber region adjacent to the source the high attenuation of particles have caused the sharp decline in the scalar flux profile. Through the ducts, for which there is no particle attenuation, the scalar flux descent is at a much lower gradient.

In order to demonstrate the capabilities of the spectral wavelets this demonstration concentrates on the scalar flux through the central region of the ducts. Figure 5.23 presents the HLW_0/HCW_0, HLW_1, HCW_1 HLW_2 and HCW_2 scalar flux solutions through the points (0,23) to (0,36). The graph shows the HLW_0/HCW_0 solution to have developed very large and dominating oscillations due to ray effects. It also reveals that many of these oscillations can be mitigated by increasing the angular expansion to either HLW_1 or HCW_1. However these higher order discretisation still form a number of smaller oscillations in their solutions. Increasing the angular expansion again to HLW_2 and HCW_2, further reductions of the oscillations are seen. However, once again visible ray effects still persist to disfigure their solutions.

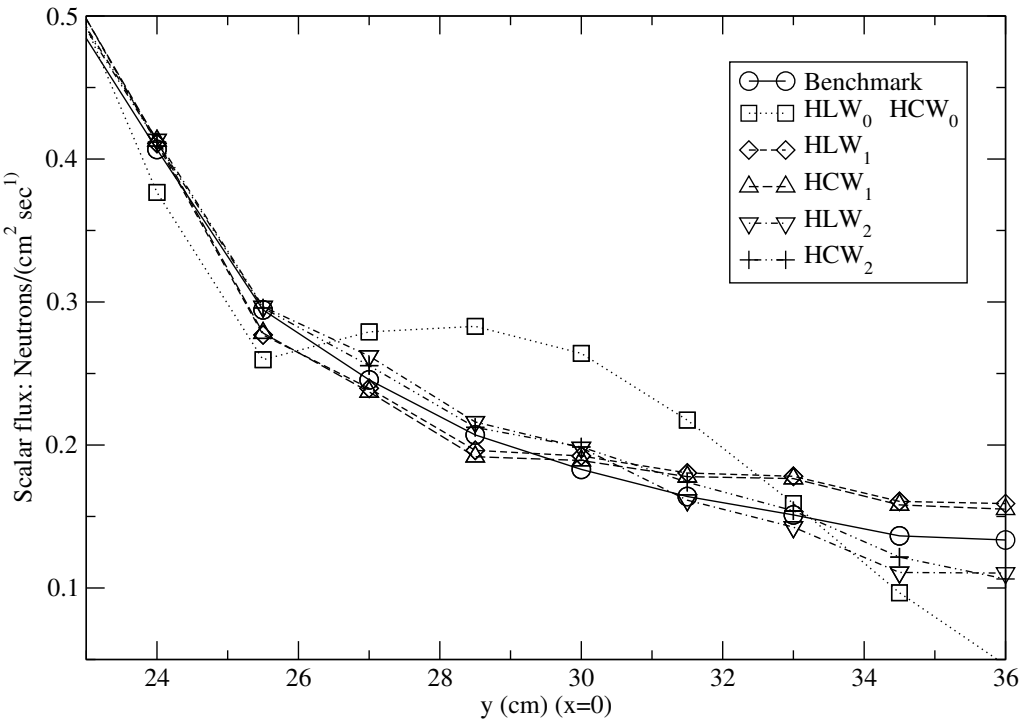

Figure 5.23: The scalar flux solutions through the positions (0,23) to (0,36) of the straight duct problem.

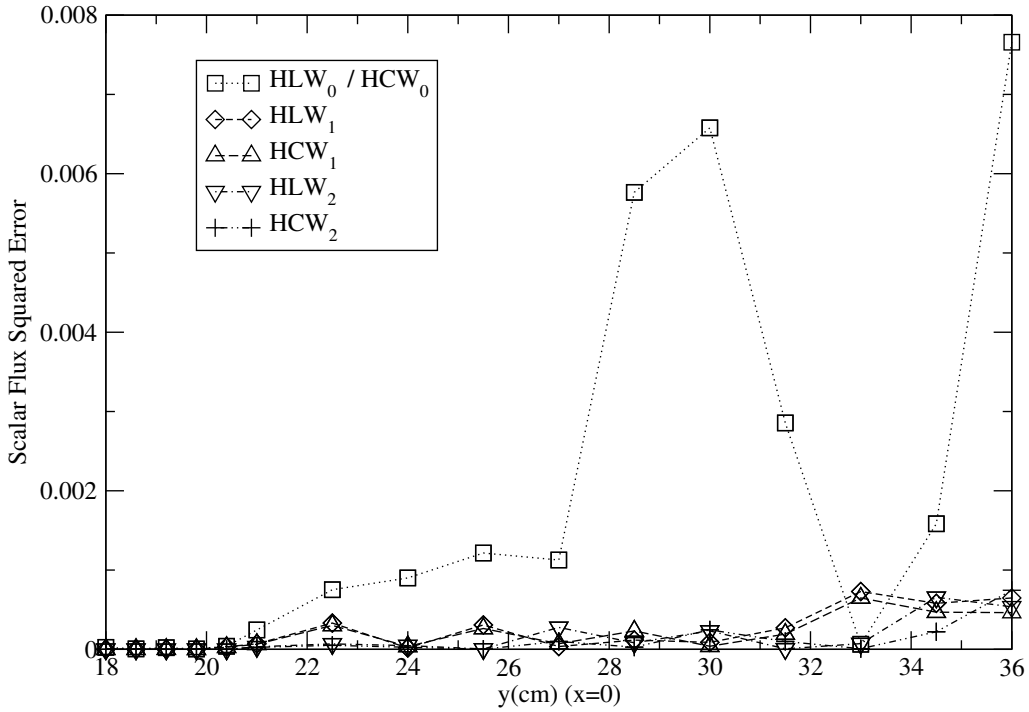

Figure 5.24: The wavelet discretisation's squared errors through the points (0,0) to (0,36).

In order to demonstrate the accuracy of the wavelet discretisations, the scalar flux squared errors are measured along the centre of the ducts. These are presented in figure 5.24. The graph reveals the extent of which oscillations have formed in the HLW_0/HCW_0 solution. Although the central region could be resolved accurately using these discretisations, the outer regions of the duct contained many large errors. The graph again demonstrates that increasing the angular resolution serves to reduces the errors significantly. In fact, the maximum squared error of HLW_0 is reduced by 90% when using HLW_1 or HCW_1.

Integrating these squared errors over the line (0,0) to (0,36) gives the error norms of the solutions through the mid sections of the duct. These values are presented in figure 5.25 which plots the error norm against the angular expansion size. We can view the graph's gradients as the representing the angular discretisation's order of accuracy in the y direction through the duct. The graph shows the spectral wavelets have the fastest rate of convergence as the error declines with the steepest gradient of -1.4 in comparison to the linear wavelets gradient of -0.9.

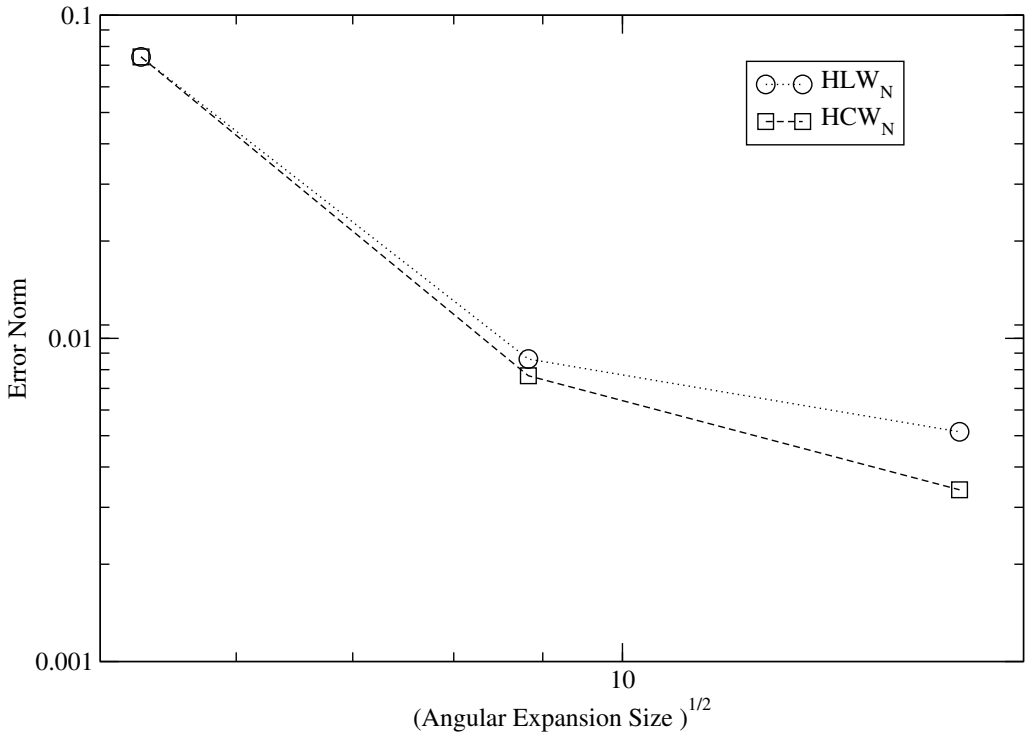

Figure 5.25: This graph shows the linear and spectral wavelet's scalar flux error norms over the positions (0,18) to (0,36).

5.5 Conclusions

In this chapter a new method for the angular discretisation of the Boltzmann transport equation has been presented. The method is a new spectral wavelet basis based on the principals of second generation wavelets. The wavelets are generated using 2 dimensional Chebyshev interpolating polynomials and approximate the angular flux over the faces the hexahedron.

A demonstration of the spectral wavelets capabilities was performed by solving three 2-D steady state neutron transport problems. The problems were designed to demonstrate the wavelet's accuracy in approximating difficult angular fluxes that may occur in materials with both optically thick and optically thin properties. Comparisons of the spectral wavelet solutions were drawn with the solutions of the hexahedral wavelets developed in chapter 4. For all numerical examples, the spatial domain was discretised using an SUPG finite element method.

The numerical examples have been able to show that the spectral wavelets are capable accurately representing the angular flux, and have therefore demonstrated their potential for use in neutron transport calculations. The method has also been shown to increase order of accuracy in comparison to the linear hexahedral wavelets developed in chapter 4. However, it should be noted that some numerical examples showed only a small increase in accuracy using the spectral wavelets. This indicates that the spectral wavelets are best suited to specific problem domains. For example, highly scattering problems where the flux varies smoothly within the angular domain would best suit the spectral method.

Combining these finding with those of the linear hexahedral wavelets in chapter 4 confirms the Chebyshev spectral wavelets to be at the very least as accurate to the established methods S_N and P_N. In addition to this, it can also be confirmed that they will be less susceptible than S_N to forming ray-effects in their solutions.

Chapter 6

SELF ADAPTIVE SPHERICAL WAVELETS FOR ANGULAR DISCRETISATIONS OF THE BOLTZMANN TRANSPORT EQUATION

Contents

This chapter presents a new adaptive angular approximation method which use spherical wavelets to discretise the Boltzmann transport equation in angle. This work builds on the work of chapters 3 and 4 in which two spherical wavelet bases were developed for the angular discretisation of the Boltzmann transport equation. A method is proposed enabling these wavelet bases to vary their angular approximation so that fine resolution in angle may be applied only to the areas of the unit sphere (representing the direction of particle travel) which are important. An error measure that operates in conjunction with the wavelet bases in order to determine this importance (and the wavelet resolution requirements on the unit sphere) is presented. A procedure by which the angular resolutions is gradually refined for steady state problems is also given.

The adaptive wavelets are applied to three test problems which demonstrate the ability of the wavelets to resolve complex fluxes with relatively few functions, and to achieve this a particular emphasis is placed on their ability to approximate particle streaming through ducts with voids. It is shown that the wavelets are capable of applying the appropriate resolution (as dictated by the error measure) to the directional component of the angular flux at all spatial positions. This method therefore offers a new and highly efficient adaptive angular approximation method.

6.1 Introduction

In simulating the transport of neutral particles the discretisation of the full 7 dimensional phase space of the Boltzmann transport equation [15] is required for deterministic solutions. However, discretisation of the 7 dimensions can quickly lead to very large sets of linear equations. This in turn can put tremendous pressure on computational resources - both in terms of time of solving the system and computational storage. For this reason it

was quickly accepted by the research community that new methods are needed in order to reduce the computational burden, for example adaptive techniques in space and angle of particle travel.

Various techniques for numerically representing the angular flux have been developed in order to reduce the costs involved in solving the discretised equations. One approach is to recast the transport equation into a single diffusion equation involving only the scalar flux, thus eliminating the angular variable completely [16, 46, 41]. The approach works well on isotropic problems, often dominated by isotropic scattering, but the method's limitations are exposed when the domain's geometry materials induce complex non-isotropic fluxes. An alternative and well used method is to form and solve the second order even-parity equation [49] in which the odd angular basis functions are eliminated. This approach avoids the restrictions of diffusion theory since angular convergence may be gained through using an infinite summation of angular basis functions like spherical harmonics (P_N). The disadvantage of this method arises when resolving particle streaming through voids or any problem with small or zero absorption cross section for which the even parity form of the transport equation breaks down. The Self Adjoint Angular Flux (SAAF) method [16] has similarities to the even parity form of the transport equation and shares these problems with it.

Adaptive angular discretisations have been developed for problems in which the distribution of angular flux varies greatly over the various regions of the system. For example, the method of composite solutions [49] eliminates the computational waste of applying high order angular resolution to all regions of a problem by having a different angular representation at each region. In many applications of interest, for example reactor systems such as the gas cooled reactors (AGRs), highly varying angular fluxes are confined to small regions of the system's geometry. It is these regions that need only be represented with a high resolution angular approximation. For the remaining regions a relatively low resolution is adequate. The method of composite solutions therefore applies a suitable P_N representation of the angular flux for each region [145, 146]. This has also been employed within the even parity equation in order to combine the efficiencies of both methods [147].

These angular adaptive methods fall into the category of isotropic methods where direc-

tional treatment may vary spatially but resolution over all angles remains uniform at any spatial point. Incorporating methods that also localise in direction, so that high resolution may be applied to only the angles containing high varying fluxes, is some what more involved [73].

This chapter presents new methods for anisotropic adaptivity in the directional treatment of the Boltzmann transport equation. To achieve this the continuous, linear, octahedral and hexahedral wavelets, chapter 3 [1] and chapter 4 [2], form the basis of the adaptive angular method. In the previous chapters, the numerical examples demonstrated the wavelet's ability to provide accurate approximations of the angular flux (without adaptivity). In many of the demonstrations they were superior to both S_N and P_N, displaying impressive capabilities in resolving void regions in situations where P_N was found to be inaccurate. They were also demonstrated to be less susceptible than S_N to ray effects [72] in the scalar flux solutions. Wavelets were used in this work because they also have hierarchical expansions and compact support. The hierarchical nature of these expansion lend themselves to the development of hierarchical (multi-grid) solvers in angle and the ability to focus resolution and computational resources on the areas of the unit sphere where they are most needed.

The wavelet expansions have been employed in the radiation transport model RADIANT [19] which employs a finite element discretisation of the space variables of the first order form Boltzmann transport equation. A Streamline Upwind Petrov Galerkin SUPG method [148, 149] was also developed within RADIANT in order to suppress spatial oscillations forming in the scalar flux solutions [88, 83, 84]. The code has been demonstrated to be highly accurate and non-oscillatory and uses a Riemann method for resolving bare surface boundary conditions [19, 89, 87]. This method maps all angular variables to a Riemann space using the eigen structure of the angular Jacobians. The incoming and outgoing information through the boundary can then be calculated and mapped back to angular variables. The importance of this is that any arbitrary angular expansion can used within RADIANT. All that is required is a re-calculation of the angular matrices with a change in the angular expansion.

The development of RADIANT has provided the possibility to investigate various angular

discretisation methods with considerable ease. To date, others research has mainly used S_N and P_N [15]. These two methods can be viewed as the two extremes of angular approximation. S_N is a first order approximation method that restricts particles to travel in a finite set of rays, directions or ordinates. Although this method leads to a simple set of hyperbolic equations coupled only through the streaming term it suffers severe oscillations in scalar flux solutions caused by ray effects [72] . The P_N approximation provides spectral accuracy and uses spherical harmonic basis functions. Although its solutions are free from ray effects (because it is rotationally invariant) it can suffer Gibbs oscillations in the angular flux. The shortcomings of these two traditionally used methods and lack of progress in developing new technologies for angular discretisation means that this important field of research is open to explore. It has already been demonstrated in chapters 3 [1] and 4 [2] that spherical wavelets are strong possible alternatives to P_N and S_N. However, in addition to this, if the application of angular adaptivity proves to be fruitful then this may lead to wavelets establishing themselves as a main tool in particle transport calculations. This can expected since the powers of wavelets have long been recognised as a tool for adaptivity to a vast range of applications.

The adaptive spherical wavelets defined by Schroder and Swelden [126] and [127], on which the ideas of our wavelets are based, have already been successfully used in a number of applications. They were able to demonstrate the wavelet's capabilities for use in compression, using topographic data relaying the depth of the earth as a function over the sphere. They were also able to demonstrate the use of wavelets in resolving bi-directional reflectance distribution functions (BRDF) and illuminations on the sphere. In these examples they were able to clarify the wavelets abilities in function reconstruction and also show the importance of the lifting scheme [91, 90]. In other words, increased accuracy was found with the lifted wavelets as opposed to their corresponding non-lifted wavelet bases. Wavelets with and without adaptivity have also been used in numerous other applications. Examples of their extensive implementation to numerical approximation of PDE's are listed in [110, 150, 111, 109].

The remainder of this chapter is organised as follows. Section 6.2.1 presents a brief description of the wavelets and their expansion series of spherical functions. This will be followed by discussion on the wavelets use in discretising the transport equation in section

6.2.2. In section 6.2.3 the SUPG angular discretised equations that are discretised in space using finite elements are formed. Then in section 6.2.4 a computationally efficient method of treating the stabilization methods is proposed. The treatment for incorporating the anisotropic adaptivity is presented in sections 6.2.5 - 6.2.6 and section 6.2.7 presents three error measures that locate areas of angle that require increased resolution. Section 6.2.8 completes the theory with a recipe for incorporating anisotropic adaptivity within an iterative solver. Section 6.3 presents calculations for three numerical examples. The numerical examples have been chosen to highlight the effectiveness of the adaptivity in resolving difficult transport problems, in particular the demonstrations concentrate on resolving neutron streaming through voids. Conclusions are contained in section 6.4.

6.2 Anisotropic Wavelets on the Sphere

6.2.1 Review of wavelet

This paper uses the continuous linear wavelets developed in [1] and [2] for the expansion of the angular flux. In order to avoid duplication, the wavelets' derivation have been omitted from here and only the details applicable to this chapter's content is reviewed. The wavelet bases represent spherical functions mapped over the octahedral and hexahedral domain presented in figures 6.1 and 6.2, respectively. Both wavelet expansions are expressed as,

$$f(\Omega) = \sum_{k \in K(0)} \lambda_{0,k}\, \varphi_{0,k} + \sum_{l=0}^{L} \sum_{k \in M(l)} \gamma_{l,k}\, \psi_{l,k}, \tag{6.1}$$

where $f(\Omega)$ represents some general spherical function. The functions $\varphi_{0,k}$ in the first term of the series are defined as the lowest order scaling functions. Together with their coefficients $\lambda_{0,k}$, they produce a basic, low resolution approximation of $f(\Omega)$ - this is employed as a starting basis for the wavelet representation. The second term of the series contains the wavelet functions $\psi_{l,k}$. These have been arranged into a double summation in order to illustrate the role they play in increasing the accuracy of the approximation of $f(\Omega)$. From level $l = 0$, wavelets are added to the expansion in a series of shells, denoted

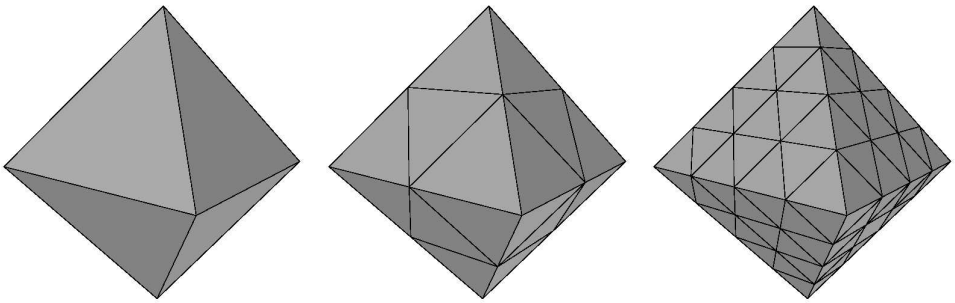

Figure 6.1: Diagram presenting the subdivision scheme employed on the faces constructing the octahedron. Left: The original faces, Middle: The first subdivision, Right: The second subdivision.

by the summation indexed l. Each shell consists of a set of wavelet functions that are denoted by the set $M(l)$, which supplements the scaling functions and wavelets of the lower ordered shells. The result is a hierarchical series of bases for the approximation of a spherical function. The hierarchical levels are generated through the inclusion of each shell of wavelet functions. The terms $\gamma_{l,k}$ in the expansion denote the expansion's coefficients corresponding to the wavelets.

The low order scaling functions are derived using standard finite element (FE) techniques. Interpolating FE linear functions centred over the nodes of the polyhedral domains are employed to generate the coarse approximating bases. The bases produce approximations with linear variation across the faces constructing the polyhedron's surface. The shells of wavelets then act upon this basic expansion, increasing the approximation's resolution in accordance with the subdivision scheme illustrated in figures 6.1 and 6.2. The subdivision scheme is a very simple procedure in which its application divides a polyhedral face into 4 new identical faces. Its relationship with the wavelets is that a series capped at level $L = j - 1$ has linear variation across the faces of the polyhedron that has undergone $j - 1$ subdivisions.

It is important to note a number of facts on the wavelet's position, size and general nature. One important property of the wavelets is that for every vertex belonging to subdivision l and not to subdivision $l - 1$, there exists a wavelet belonging to the shell $M(l)$ that

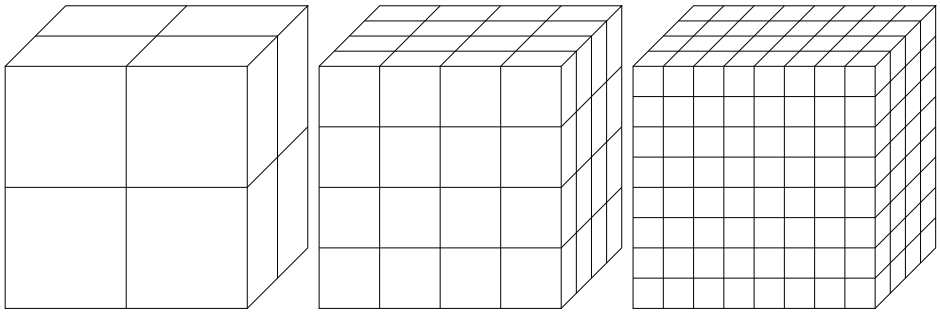

Figure 6.2: Diagram presenting the subdivision scheme employed on the faces constructing the hexahedron. Left: The original faces, Middle: The first subdivision, Right: The second subdivision.

is associated with that vertex. The vertices are defined as the points on the polyhedral surface generated by the subdivision scheme. Stating that a wavelet is associated with a vertex means that the wavelet is positioned over that point and is non-zero in the local vicinity of that point. That is, they have compact support which implies that the area in which the wavelet is non zero in value is restricted to a strict subset of their domain. The nature of the wavelets are that as their order in the expansion shells increase, the supports of the functions reduce in size. The implication of this is that the wavelets' information becomes heavily concentrated about their vertices as the order of the wavelets increase. Some examples of the octahedral wavelets are presented in figure 6.3.

It should be noted here that the subdivision scheme ensures all vertices are nested, that is that the vertices of subdivision $l-1$ are also included in subdivision l, for all values of l. It is also important to observe that the new vertices generated on an application of the subdivision scheme are positioned mid way between all the existing connected vertices - figure 6.4 presents an illustration for various sub division levels. Not only does this ensure that the wavelets within each shell are distributed homogeneously over the polyhedral domain, but also every wavelet (and scaling function) has close neighbouring wavelets belonging to the next level of the hierarchy - also see figure 6.5.

For the present purposes the important wavelet properties are:
1) Wavelets have hierarchical expansions.
2) Wavelets have compact support that reduce in size around their collocation point (node)

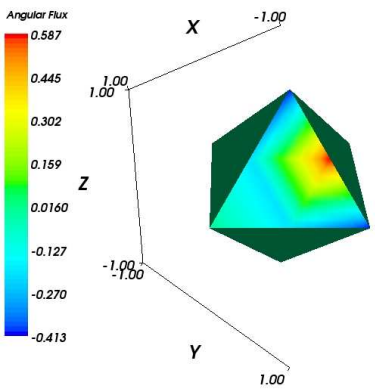

(a) Level 1 wavelet

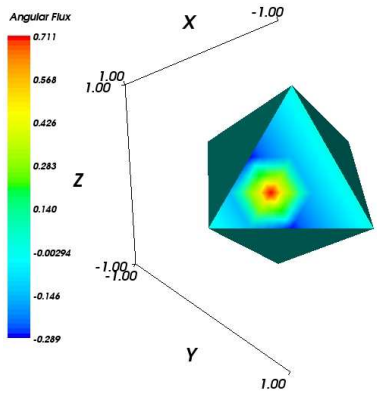

(b) Level 2 wavelet

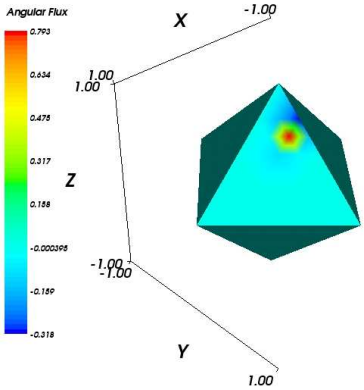

(a) Level 3 wavelet

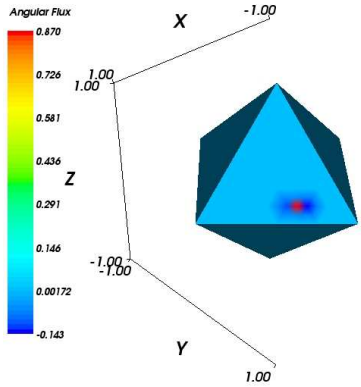

(b) Level 4 wavelet

Figure 6.3: Diagram displaying some wavelet functions over the octahedral domain. As the order of the wavelet increase their support become highly concentrated about their respective point.

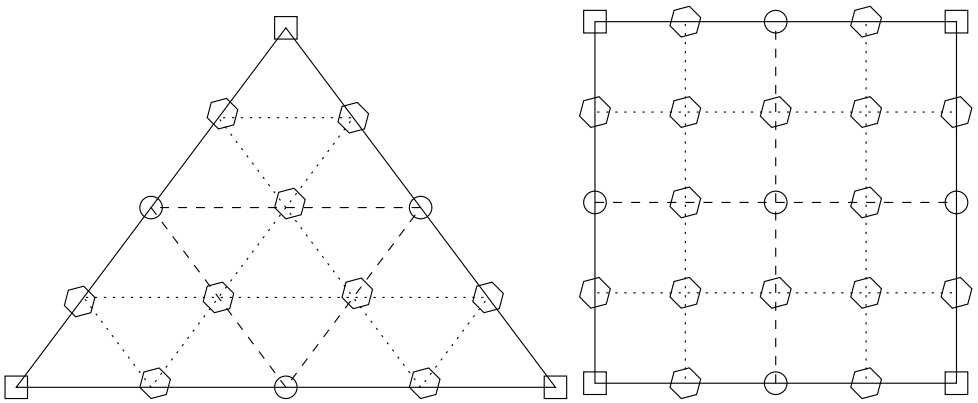

Figure 6.4: Diagram displaying a face of the octahedron (left) and hexahedron (right). The squares represent the position of the lowest order scaling functions. The circles represent the positions of the first level of wavelets and the hexagons represent the positions of the second level wavelets. It is shown that for every wavelet, there exists a set of neighbouring wavelets belonging to the next higher level. The solid, dashed and dotted lines represent the face partitioning on the zeroth, first and second subdivision, respectively.

as the expansion increases.
3) Wavelets of all levels have 'close by' neighbouring wavelets associated with the next hierarchical level - these are the new mid-point vertices generated in the subdivision - see figure 6.4 and 6.5.

6.2.2 The angularly discretised Boltzmann transport equation

In order to resolve the directional variable of the BTE employing an adaptive angular technique the equation's directional dependence is first discretised using a standard spherical wavelet method. A wavelet angular discretisation of the Boltzmann transport equation requires representing the angular variable of ψ by the series (6.1). That is the angular flux ψ is approximated by,

$$\psi(r,\hat{\Omega},t) \approx \sum_{j=1}^{\mathcal{M}} \Psi_j(r,t)\mathcal{G}_j(\hat{\Omega}), \tag{6.2}$$

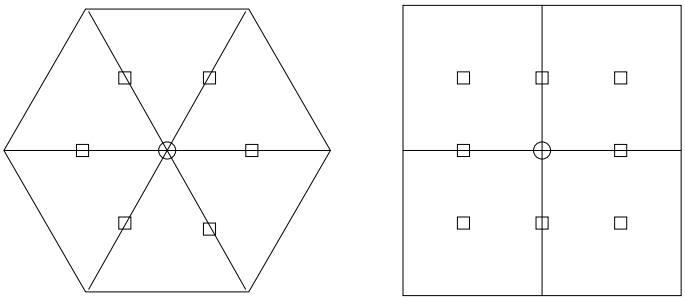

Figure 6.5: Diagrams displaying the the neighbouring functions for each wavelet. The ◯ represents the wavelet of level j under consideration. The □ represents the wavelets of level $j+1$ that are included into the expansion should the approximation around ◯ require further resolution.

where $\mathcal{G}_j, \forall j \in \{1, \ldots \mathcal{M}\}$, represent the wavelets basis functions and Ψ_j are their corresponding coefficients. Note that a suitable reordering of (6.1) is used to reduce the series into a single summation. The approximation is then inserted into the first order Boltzmann transport equation [15] and the standard Galerkin method is applied resulting in the angularly discretised transport equation.

$$\frac{1}{v}A\frac{\partial\Psi}{\partial t} + A_x\frac{\partial\Psi}{\partial x} + A_y\frac{\partial\Psi}{\partial y} + A_z\frac{\partial\Psi}{\partial z} + \mathcal{H}\Psi = \mathcal{S} \tag{6.3}$$

The matrix A is the $\mathcal{M} \times \mathcal{M}$ angular mass matrix. The matrices A_x, A_y and A_z are the $\mathcal{M} \times \mathcal{M}$ angular Jacobians. The $\mathcal{M} \times \mathcal{M}$ matrix H represents the scattering and removal operator and $\mathcal{S}$ is the external source vector of size $\mathcal{M}$. The vector Ψ represents the $\mathcal{M}$ coefficients of the wavelet expansion in (6.2). A complete derivation of this angularly discretised equation, together with the definition of the matrices, can be found in [1] and [2] and chapter 2.

6.2.3 Streamline Upwind Petrov Galerkin spatial discretisation

In order to resolve the spatial dimensions of the Boltzmann transport equation a finite element method is used. The approach uses a weighted residual method in order to generate the weak form of equations that are solved in terms of the spatial discretisation's

unknown coefficients. A Streamline Upwind Petrov Galerkin method is also used, the basic mechanics of this being that a contribution of the convection term is added to the weighting of the angularly discretised equations. This induces more diffusion across the spatial dimensions which in turn helps to prevent Gibbs oscillations forming in the solution. The Streamline Upwind Petrov Galerkin formulation of the angular discretised equations (6.3) is given by,

$$(I - A \cdot \nabla P)\mathcal{R}(\Psi(r,t)) = 0, \tag{6.4}$$

where

$$\mathcal{R}(\Psi(r,t)) = \frac{1}{v}A\frac{\partial \Psi(r,t)}{\partial t} + \boldsymbol{A} \cdot \nabla \Psi(r,t) + \mathcal{H}\Psi(r,t) - \mathcal{S}(r,t). \tag{6.5}$$

The term $\boldsymbol{A}$ denotes the vector of angular Jacobians and $\nabla = (\frac{\partial}{\partial x}I, \frac{\partial}{\partial x}I, \frac{\partial}{\partial y}I)$. P is the $\mathcal{M} \times \mathcal{M}$ stabilisation matrix. Its coefficients determine the dissipative properties of the discretisation method. Various formulations for this matrix exist of which some are discussed in [84]. The stabilising term used in this chapter is treated in the next section.

In this work time is neglected and the spatial variables of equation (6.4) are approximated using finite element interpolating functions [49]. The angular discretised flux Ψ, scattering and removal operator matrix $\mathcal{H}$ and source terms $\mathcal{S}$ are spatially represented by,

$$\Psi(r) \approx \boldsymbol{\Psi}(r) = \sum_{j=1}^{\mathcal{N}} \boldsymbol{N}_j(r)\boldsymbol{\Psi}_j, \tag{6.6}$$

$$\mathcal{H}(r) \approx H(r) = \sum_{j=1}^{\mathcal{N}_\mathcal{H}} \boldsymbol{N}_{Hj}(r)H_j, \tag{6.7}$$

and

$$\mathcal{S}(r) \approx S(r) = \sum_{j=1}^{\mathcal{N}_\mathcal{S}} \boldsymbol{N}_{Sj}(r)S_j, \tag{6.8}$$

respectively. The matrices $\boldsymbol{N}_j(r), \boldsymbol{N}_{Hj}(r)$ and $\boldsymbol{N}_{Sj}(r)$ are $\mathcal{M} \times \mathcal{M}$ diagonal matrices containing the finite element function N_j, N_{Hj} and N_{Sj}, respectively, along their diagonal. The spatial approximation is performed by the set of finite element functions $\boldsymbol{N}_j$, for $j \in \{1, 2, \ldots, \mathcal{N}\}$. The term $\boldsymbol{\Psi}_j$ is a vector of size $\mathcal{M}$ that represents the angular moments at node j of the finite element mesh. The term H_j is an $\mathcal{M} \times \mathcal{M}$ matrix denoting the angular discretised scattering/removal operator at node j of the finite element mesh. Finally, the

vector $S_j(t)$ is of size $\mathcal{M}$ and contains the angularly discretised source terms associated with node j of the finite element mesh. In this chapter the condition $N_j = N_{Hj}$ = N_{Sj} is assumed for $j \in \{1, 2, \ldots, \mathcal{N}\}$.

To obtain the weak FE equations, equation (6.4) is pre-multiplied by the $\mathcal{M} \times \mathcal{M}$ matrix $\boldsymbol{N}_i(r)$. This is then integrated over space resulting in:

$$\begin{aligned} -\int_V (\boldsymbol{A} \cdot \nabla \boldsymbol{N}_i(r))\Psi(r)dV &+ \int_V \boldsymbol{N}_i(r)(H(r)\Psi(r) - S(r))dV \\ +\int_V \boldsymbol{A} \cdot \nabla \boldsymbol{N}_i(r) P\mathcal{R}(\Psi(r))dV &+ \int_\Gamma \boldsymbol{N}_i(r)(\boldsymbol{A} \cdot \boldsymbol{n})(\Psi(r))d\Gamma \\ -\int_\Gamma \boldsymbol{N}_i(r)(\boldsymbol{A} \cdot \boldsymbol{n}) P\mathcal{R}(\Psi(r))d\Gamma &= 0, \qquad \forall\, i \in \{1, \ldots \mathcal{M}\} \end{aligned} \tag{6.9}$$

Here the surface integrals are a product of an application of Green's theorem and $\boldsymbol{n}$ denotes the direction normal to the surface of the problem's space boundary Γ. The identity matrix of size $\mathcal{M} \times \mathcal{M}$ is represented by I. The terms $\boldsymbol{A} \cdot \boldsymbol{n} = n_x A_x + n_y A_y + n_z A_z$ and $\boldsymbol{A}.\nabla = A_x \frac{\partial}{\partial x} + A_y \frac{\partial}{\partial y} + A_z \frac{\partial}{\partial z}$. In this formulation the last term on the left hand side of equation (6.9) is assumed to be zero.

6.2.4 Definition and approximation of P

Incorporating the correct diffusion within the solution's streamline direction is critical for the mitigation of Gibbs oscillations manifesting over regions with large flux gradients. A number of stabilising methods defining the term P have already been developed that generally perform well when resolving transport problems with specific material cross sections. This chapter employs the method proposed in [19] and [84] for which the stabilisation matrix combines two popular approaches in an attempt to extract the best qualities from each. The formal definition of P used in this work is given by,

$$P = f_{min}\Big(\frac{h}{2}(f_2(A)^{-1}, H^{-1})^T\Big) \tag{6.10}$$

for some vector of matrices $B = (B_1, B_2, \ldots, B_l)$. The term f_p is defined by,

$$f_p(B) = ||B_1|^p + |B_2|^p + \ldots + |B_l||^{\frac{1}{p}}. \tag{6.11}$$

The operator $|B_i|^p$ is defined as,

$$|B_i|^p = L_{B_i}^{-1} |\Lambda_{B_i}|^p R_{B_i}^{-1} \tag{6.12}$$

for which L_{B_i} and R_{B_i} denote the matrices formed from the augmentation of the left and right eigen-vectors of B_i, respectively, and Λ_{B_i} is the diagonal matrix containing the eigen-values of B_i. h in 6.10 denotes a measure of the elements length, see [84].

The definition of P in its current form (6.10) gives rise to two issues with regards to its practical implementation. First there are a number of mappings involved, each requiring matrix matrix multiplications using the eigen structure matrices of the angular Jacobians. These computations will need to be performed for each spatial node of the finite element mesh since the inclusion of the H and $\frac{h}{2}$ terms makes P spatially variable. These mappings are computationally expensive when using large angular expansions. It is therefore necessary to develop an approximation of P in order to reduce this computational burden. The second issue is with regard to the computation and sparsity of the matrices generated by the multiplications of the angular Jacobians and P. That is the matrices $A_x P A_x, A_x P A_y, A_x P A_z, \ldots$, generated in the third term of (6.9). The computation and application of these matrices is highly expensive if they are allowed to become dense and it is highly possible that this will occur with the current P. In order to keep the matrices to a reasonable level of sparseness some restrictions on the approximation of P must be imposed.

To address these computational concerns the following assumptions are implemented. First $f_2(A)^{-1} = A^{-1}$ where A is the angular mass matrix. This is exact in 3-D calculations and will also give close approximations in 2-D. By using this the first set of mappings involved in calculating $f_2(A)^{-1}$ are eliminated. In order to address the second issue, that is in keeping the matrices in the third term of (6.9) sparse, P should also be sparse, see below.

Diagonal P

One method is to impose the restriction on P that it is diagonal. This can be ensured by using diagonal approximations to $f_2(A)^{-1}$ and H^{-1} in equation (6.10). To implement this an approximation of the matrices are used. A diagonal approximation of the matrix H is

formed through a row summation of it absolute values. The approximation is denoted by $H^{'}$ and is given by,

$$H^{'}_{i,j} = \begin{cases} \sum_k |H_{i,k}| & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

In order to obtain a diagonal approximation of A^{-1} the inverse of the diagonal elements of A are used:

$$A^{'}_{i,j} = \begin{cases} \frac{1}{A_{i,j}} & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

For orthonormal angular discretisations, for example P_N , S_N and Walsh functions [77, 78], this is exact since the angular mass matrix A is already diagonal. Since the wavelets are not orthonormal the above definition will be an approximation. However, the approximation is good due to the diagonal dominance of the angular mass matrix A.

Combining the above amendments, the approximation to P is now defined as,

$$P^{'}_{i,j} = \begin{cases} min((H^{'}_{i,j})^{-1}, \frac{h}{2} A^{'}_{i,j}) & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Expansion of P

An alternative route for generating a sparse approximation of P is to expand the inverse of the angular mass matrix by the following sum of matrices,

$$A^{-1} = A^{'}_{m} \approx (\alpha_0 I + \alpha_1 V + \alpha_2 V^2 + \ldots + \alpha_{m-1} V^{m-1}) V_1^{-1}, \tag{6.13}$$

where $V = I - V_1^{-1} A$ and V_1 and V_2 are generated through some splitting of the matrix $A = V_1 + V_2$. A possible definition of the constants $\alpha_{\cdot}$ is presented in [151]. The appealing aspect of this approach is that the low order matrices can be made to retain a high order of sparseness and so by truncating the summation to just the first two or three terms, a sparse and reasonably accurate approximation of A^{-1} is generated.

6.2.5 Application of adaptivity

Developing methods of anisotropic adaptivity using arbitrary angular expansions would normally be plagued with numerous complications. The greatest problem being that a varying angular discretisation at each spatial position would destroy the structure of the angular discretised transport equation (6.3). However, the use of hierarchical expansions provides a possible route for avoiding this situation. The procedure considers the employment of a full, high order wavelet expansion for discretising the transport equation that produces a system of equations in the form of (6.3). Then, at each spatial point, the wavelets that do not contribute to the angular flux approximation are cancelled out by assuming the coefficient corresponding to the redundant wavelet has the value zero - this would indeed be true if the wavelet has no contribution. The consequence of this is that the wavelet is removed from consideration resulting in a reduced angular expansion. This procedure therefore gives us the required freedom of having varying expansion of angular basis functions for every spatial position, whilst at the same time keeping the mathematics watertight by considering the complete, spatially non-varying angular approximation. The full potential of the wavelets can now be realised through their property of having compact support as this presents a straight forward method for implementing anisotropic adaptivity. Compactly supported functions allow the application of high resolution to only the regions on the unit sphere where the angular flux variation is important. For example, if the angular flux is non-zero in a small area on the sphere then only the wavelets with supports covering this region of angle are required for an accurate representation.

6.2.6 Wavelet selection process

The wavelet bases naturally lead to a method that determines the regions on the unit sphere in which the angular flux is under resolved when using some truncated wavelet expansion. Furthermore they also present a simple but effective method that selects the wavelets that will increase resolution. The directions in which the angular flux is under resolved can determined as the regions surrounding wavelets with large expansion coefficients Ψ_j in equation (6.2). The reason for this can be explained with the aid of the diagrams

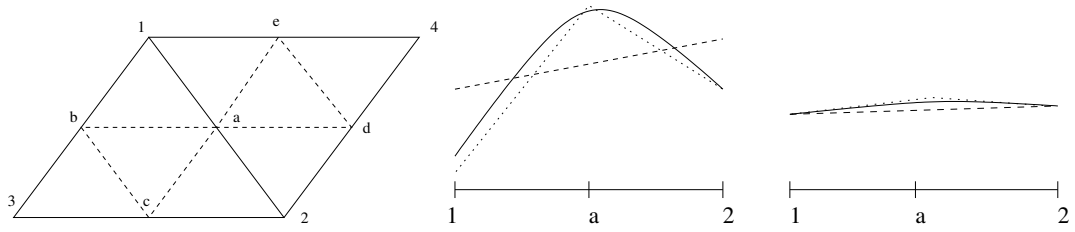

Figure 6.6: Left: A diagram displaying a subset of the octahedral face. The points labelled 1-4 represent wavelets of subdivision j. The points labelled a-e represent wavelets of level $j+1$. Middle and right: Diagrams displaying how the wavelet of position a operates to increase the accuracy of a function approximation around its associated vertex. The solid line represents the function being approximated. The dashed line represents the approximation of subdivision j. The dotted line represents the approximation at level $j+1$. The line plots are drawn along the edge 1-2.

presented in figure 6.6 representing a subset of the octahedral domain. The diagram on the left presents two levels of the hierarchical subdivision. The coarser level is indicated by the bold line triangles and associated with it are the wavelets $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ and $\mathcal{G}_4$. The higher hierarchical level is shown by the thin line triangles and associated with it are the wavelets $\mathcal{G}_a, \mathcal{G}_b, \mathcal{G}_c, \mathcal{G}_d$ and $\mathcal{G}_e$. The coefficients corresponding to the higher order wavelets indicate whether the lower ordered wavelets were sufficient in approximating the angular flux around the area of its support. In particular, when increasing the approximation, the coefficients of the higher order wavelets determine the difference in the angular flux approximation over their corresponding vertices. This is illustrated by the two examples presented in figure 6.6. Both examples show the flux as a line through the edge joining vertices 1 and 2 and show the low and high resolution approximation to some curve. The first example shows the poorly performing low resolution approximation requiring a large contribution from the wavelet $\mathcal{G}_a$ in order to correct the large error across its collocation point. This in turn requires the coefficient corresponding to the wavelet $\mathcal{G}_a$ to be relatively large. However, the low resolution approximation of the second example was sufficient to approximate the curve. This therefore required little or no contribution from the wavelet $\mathcal{G}_a$ and thus leads to a small coefficient Ψ_a.

6.2.7 Error measures

This section develops error measures based on the magnitude of the wavelet coefficients in order to locate the areas of angle in which the angular flux is under resolved. On detecting a large error around a wavelet collocation point the procedure adds the neighbouring wavelets of the next hierarchical level to the expansion. These wavelets, discussed in summary point 3 of section 6.2.1, are presented in figure 6.5. These wavelets will serve the purpose of increasing the resolution of the important areas on the unit sphere and also indicate whether further resolution is required. This process is repeated until all wavelets necessary for an accurate approximation over all angles are included.

Here three error measures based on the magnitudes of the wavelet's coefficients are proposed.

Error measure 1 is a measure of the variation introduced by a wavelet over its associated vertex. Its value takes the wavelet's height over the collocation point multiplied by its coefficient. A large error is then registered over the collocation point if this value exceeds a predetermined tolerance τ. This is explicitly written as,

$$\mathcal{G}_i(\Omega_i)\Psi_i \geq \tau, \tag{6.14}$$

where Ω_i is the angular position of the node associated with the wavelet $\mathcal{G}_i$.

Error measure 2 analyses the scalar flux contribution from a wavelet. The error measure is given by the wavelet's height multiplied the coefficient and the area covered by the function's support. This is aimed at providing a fairer reflection of the wavelet's contribution toward the scalar flux, as low order wavelets (with large areas of support) can make large contributions with small coefficients. The error measure is given as,

$$\mathcal{G}_i(\Omega_i)\Psi_i\mathcal{A}_i \geq \tau, \tag{6.15}$$

where $\mathcal{A}_i$ represents the area of the wavelet's support.

$$\mathcal{A}_i = \int_{G_i \neq 0} d\Omega \tag{6.16}$$

This value can be approximated by summing the areas of the triangles spanned by the wavelets.

Error measure 3 normalises the second error measure with respect to the scalar flux. This may be useful as the value will represent the wavelets contribution with respect to the actual amount of flux at a spatial point. This measure is given by,

$$\frac{\mathcal{G}_i(\Omega_i)\psi_i\mathcal{A}_i}{\max\{\phi,\epsilon\}} \geq \tau, \tag{6.17}$$

where ϕ represents the scalar flux, $\phi = \int_\Omega \psi d\Omega$, and ϵ represents a small number greater than zero. The inclusion of ϵ ensures that unnecessary resolution is not placed in regions containing negligible scalar fluxes.

6.2.8 Self adapting wavelets within an iterative solver

Here the anisotropic adaptivity is incorporated to work within a general iterative solver, for example a block Gauss-Seidel [152] or some Krylov subspace method such as GMRES [153] or FGMES [154]. The mechanics of the solver works by improving a current approximation through finding a better approximation by searching along a specific search direction. The search process is iterated until the exact solution is found or the approximation is within some tolerance of the true solution. To perform adaptivity in angle the following is used:

1) The discretised equations (6.3) are formed with some maximum angular expansion.
2) Initial starting conditions are that only the lowest order scaling functions are included in the angular expansion for all spatial points.
3) During each iteration of the solver, the coefficients are checked and, if they meet the criteria for adaptivity, the neighbouring wavelets are added to the angular expansion.
4) The solver is allowed to iterate until a convergence has been reached.

6.3 Numerical Examples

In this section 3 steady-state problems are investigated in order to illustrate the capabilities of the adaptive wavelet discretisation method. In resolving these numerical examples

error measure 2 is employed using a range of different thresholds τ. The aims are to illustrate the adaptivity working in conjunction with the error measure and to investigate the tolerance value's effects on the numerical approximations. In order to give a fair view of the robustness of the adaptive wavelets, the problems have been selected to expose poorly performing angular discretisations. In particular, much emphasis has been placed towards on resolving angular fluxes within voids.

In order to give some indication to how the adaptive methods increase the efficiency of solving the following numerical examples, the following assumptions have been made. The calculation times for the adaptive and non adaptive methods are calculated only in terms of the angular expansion size (as the spatial discretisation remains the same for each problem its influence on the solving time also remains the same and so can be neglected). For the non-adaptive discretisations, the calculation times will be dependent on the angular expansion size $\mathcal{M}$ that is defined in equation 6.2. For the adaptive methods, the calculation times are based on the average number of angular basis functions that were used on the spatial nodes of the finite element mesh, this value will be denoted by $\mathcal{M}_{adp}$.

For the numerical example in this chapter, the preconditioned GMRES algorithm has been used to solve the linear systems generated from the discretised equations. The time complexity for solving the linear systems are therefore dominated by the time to compute a matrix vector multiplication and the application of a preconditioner. The preconditioner is a block SSOR sweep over the finite element spatial nodes. This involves building and solving a matrix of dimension $\mathcal{M} \times \mathcal{M}$ or $\mathcal{M}_{adp} \times \mathcal{M}_{adp}$ to resolve the angular moments at each node of the spatial mesh. Therefore, the time to compute the matrix vector multiplication and preconditioner will essentially depend on these block matrices. In particular, the number of computer operations will depend on the sparsity of their non-zero elements. We can therefore make upper and lower bounds on the time complexity for these operations. In the best case, where the block matrices are very sparse, the time complexity is of order $\emptyset(\mathcal{M})$ (or $\emptyset(\mathcal{M}_{adp})$). In the worst case, where the block matrices have a high proportion of non-zero elements, the time complexity is of the order $\emptyset(\mathcal{M}^2)$ (or $\emptyset(\mathcal{M}_{adp}^2)$). In the following numerical examples, the relative approximate running times of the adaptive wavelet methods are calculated against the approximate calculation times of the full wavelet expansion using these upper and lower bounds.

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.0
2	0.0	0.5	0.0

Table 6.1: Source and material properties of the Box Source Problem 1.

The following numerical examples have all been restricted to time independent 1 group problems. The wavelet calculations have been spatially solved using RADIANT. The results of these numerical examples are compared with two established radiation methods. The first method, embedded in the EVENT code [20], is based on the variational formulation of the even-parity form of the transport equation, employing continuous finite elements and spherical harmonic basis functions to discretise the space and angle variables [129]. The second method embedded in the DG-DO code solves the first order Boltzmann transport equation using discontinuous finite elements (a discontinuous Galerkin method) and discrete ordinates to discretise spatial and angle variables [21].

6.3.1 Pure absorbing box source problem

This example is designed to demonstrate the adaptive linear wavelets capabilities in resolving a simple 2-D problem. The adaptive method is compared for various tolerance settings in order to gain an insight of the method's sensitivity of adapting, with respect to this value.

This example uses the 2-D box source problem presented in figure 6.7. The corresponding source and material properties are listed in table 6.1. The domain is a 3cm $\times$ 3cm square with a constant absorbing cross-section $\sigma_a = 0.5\mathrm{cm}^{-1}$. An isotropic source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ occupies the central square region of size 0.5cm $\times$ 0.5cm. Vacuum boundaries are prescribed along all sides of the problem domain. For all calculations, the 2-D box source problem was solved using rectangular shaped bi-linear finite elements. This calculation has 225 quadrilateral elements, ranging from 0.1$\times$0.1 cm to 0.2$\times$0.2 cm in size, and 256 nodes.

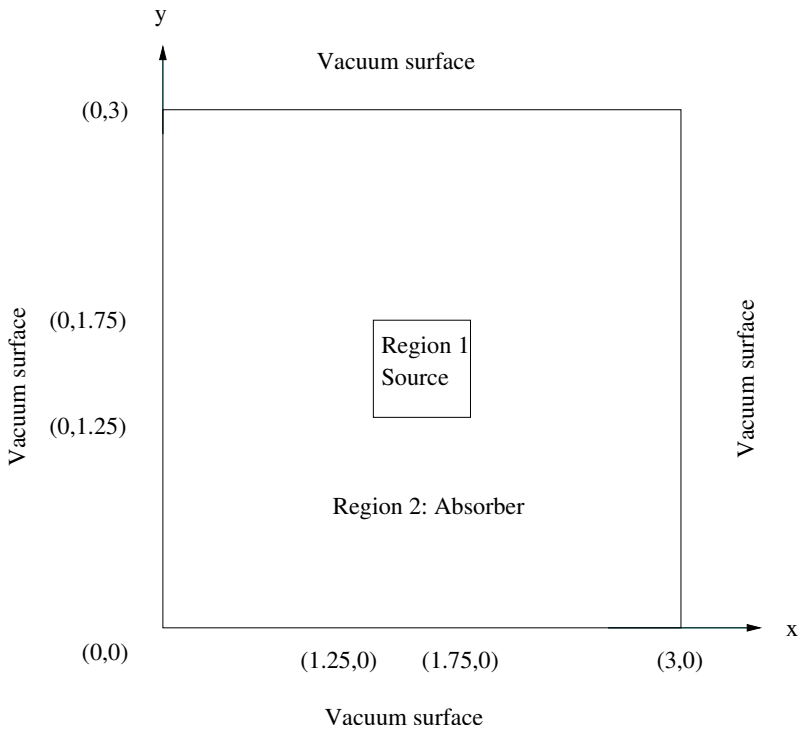

Figure 6.7: Diagram of the 2-D Box Source Problem.

The numerical solutions were provided by the adaptive octahedral wavelets and allowed a maximum of 612 angular basis functions per spatial node (that is 4 layers of wavelets plus the scaling functions, this is denoted by LW_4). Error measure 2 was employed using the three tolerance values 0.01, 0.005 and 0.001. The solution to this problem using the lowest tolerance is shown in figure 6.8. It is possible to confirm that the scalar flux profile is in very close agreement with the solutions benchmarked in [1] and [2]. The solution peaks over the central region reaching approximately 3.4 particles.$\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The profile then falls away sharply from the centre as particles attenuate within the absorbing material.

Figure 6.9 shows three graphs displaying the number of wavelets used in the angular expansions over the finite element nodes of the spatial discretisation. The solutions confirm that a relaxation in the error measure leads to significant increase in the inclusion of the wavelet functions. The results from the tolerance $\tau = 0.01$ required on average 70 wavelet functions per spatial node. This value increased to 122 wavelets in the calculation using $\tau = 0.005$ and jumped significantly to 246 wavelets for the lowest tolerance calculation

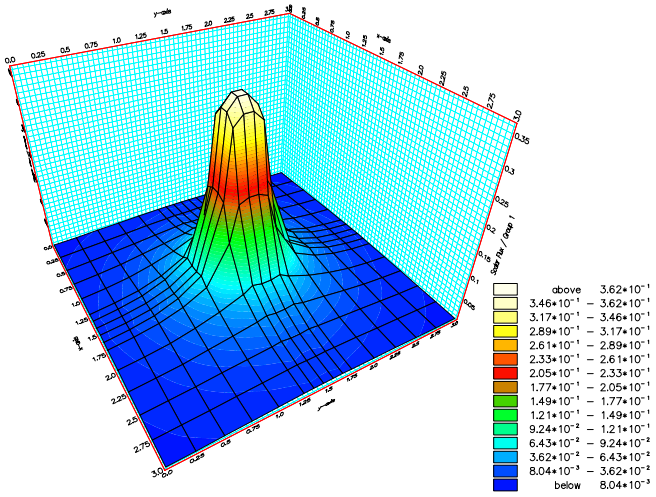

Figure 6.8: Scalar flux profile of the adaptive solution using τ = 0.001.

$\tau = 0.001$. Table 6.2 lists the relative computation times for these three calculations using the best and worst bounds described at the start of this section. The figures show that the adaptive technique can reduce the computation cost from between 1.3% and 40.1% depending on the tolerance setting used.

It is however important to observe that the low and more accurate threshold τ produces a pattern in the expansion size that is uniform with respect to the distance from the problem's centre. This would appear correct since fluxes at equal distance from the problem's centre will be of similar distribution and, therefore, require a similar number of angular basis functions. This pattern is not replicated in the calculation using the threshold $\tau = 0.01$. Here this tolerance appears to have restricted many of the wavelets from entering the calculation - especially in the diagonal directions with respect to the source. However, the threshold $\tau = 0.005$ appears to be a compromise between the two extreme values.

All three results shown in figure 6.9 are in agreement in that the angular flux within the central source region requires the largest wavelet expansions. This is due to complex angular distributions covering the whole sphere being generated by the 2-D source embedded within the weakly absorbing media. These type of geometries typically generate high maximum around the poles of the sphere and require high resolution near the poles

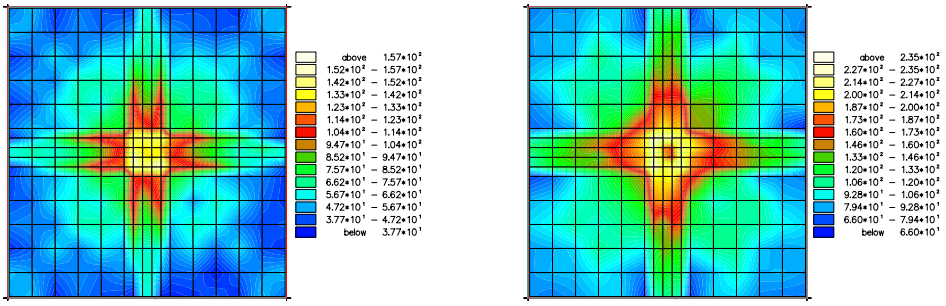

(a) τ = 0.01 (b) τ = 0.01

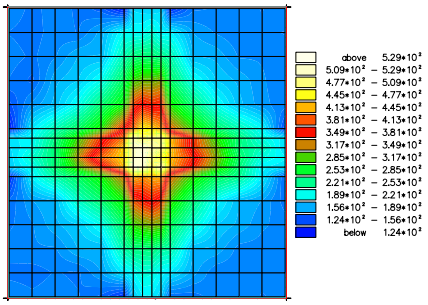

(c) τ = 0.001

Figure 6.9: Diagram plotting the number of angular basis functions used over the spatial grid. Top left: τ = 0.01, Top right: τ = 0.005, bottom: τ = 0.001

tolerance τ	Angular basis functions($\mathcal{M}_{adp}$)	Best relative time	Worst relative time
0.01	70	1.3 %	11.4 %
0.005	122	3.9 %	19.9 %
0.001	246	16.1 %	40.1 %

Table 6.2: The table shows the average number of angular basis functions used for resolving problem 1. Included are the best and worst relative computation times to solve the adaptive equations - in comparison to the full wavelet expansion.

in order to resolve the flux fully. Along the boundary there is a significant drop in the expansion sizes for all three results. This gives a strong indication that adaptivity is working correctly as these positions contain highly directional fluxes - thus making a large proportion of wavelets redundant. On average, the increasing tolerance values require 45, 90 and 160 wavelets in order to resolve the angular flux positioned along the boundary. This shows the angular expansion size for the $\tau = 0.01$ calculation to be notably smaller than the other two. However this vastly reduced expansion may also result in an under resolved approximation of the angular flux. This in fact turns out to be confirmed by the graph contained in figure 6.10. The graph shows the scalar flux profile along the bottom boundary from position (0,0) to position (3,0). Large oscillations have distorted the solution of the $\tau = 0.01$ calculation due to the lack of resolution given to the approximation of the angular flux. However the figure does show that decreasing the tolerance to $\tau = 0.005$ was sufficient to eliminate these oscillations across the boundary interface. The graph also reveals the solutions of the 0.001 and 0.005 calculations to have similar profiles, indicating a level of convergence in terms of tolerance value has been found at $\tau = 0.005$. This is confirmed by the benchmark solution provided by the DG-DO method shown in figure 6.10.

Figure 6.11 presents three angular flux profiles obtained from the solutions of the various tolerance values sampled on the boundary at position (1.35,0). The figure also includes the wavelet position plots corresponding to these three angular flux results. In these plots points are placed over the collocation points of the wavelets that are included in the expansion of the angular flux in order to help illustrate the positions on the sphere where high resolution has been applied. The three profiles show the angular flux is concentrated along a vertical strip directing particles away from the source region in the negative y direction. The profiles also show a higher concentration of particles are directed in the negative x direction which conforms with the slight off centre position of the sampled point.

The three results agree with the findings already observed with this numerical example. By comparison of all angular flux results it is clear that the tolerance $\tau = 0.01$ is too high for a good reconstruction of the angular flux. The angular flux profile obtained using this tolerance shows poor representation of the peaked strip which possesses a number of large

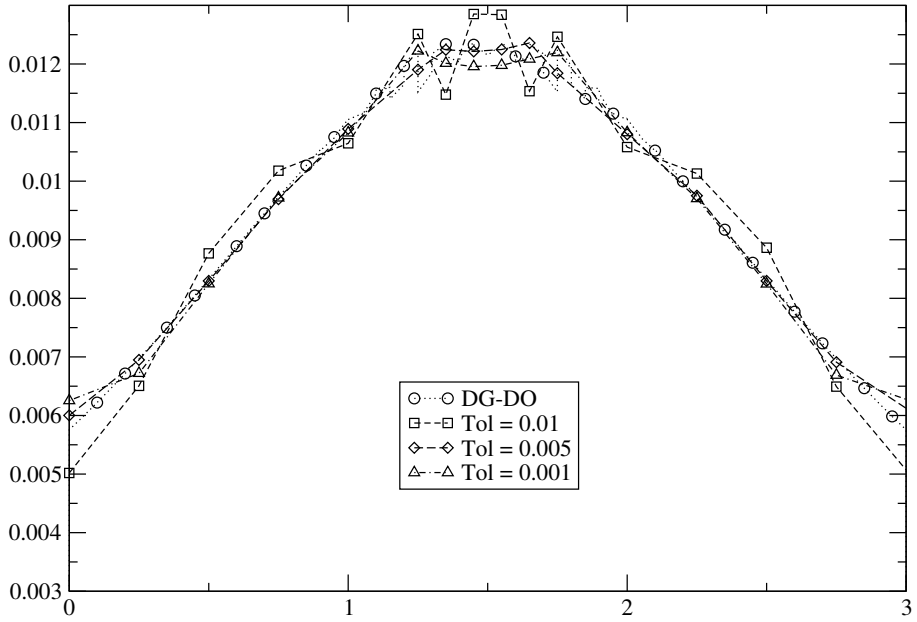

Figure 6.10: Scalar flux profiles of the solutions obtained from DG-DO and adaptive wavelets using tolerance $\tau = 0.01, 0.005$ and 0.001. The line plots were sampled along the box source problem's bottom boundary (0,0)-(3,0).

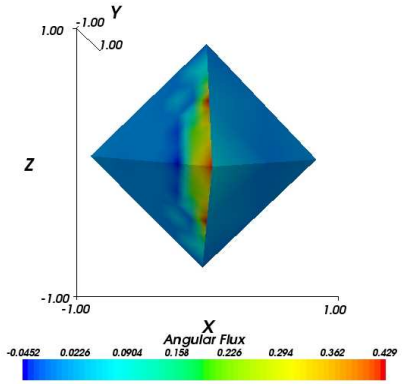

(a) Angular flux: τ = 0.01

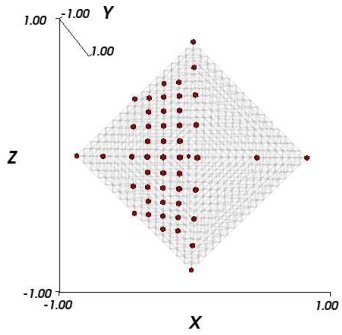

(b) wavelets: τ = 0.01 (38)

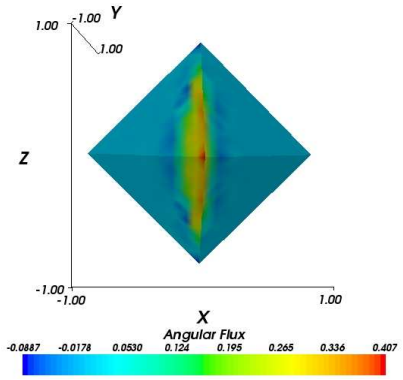

(c) Angular flux: τ = 0.005

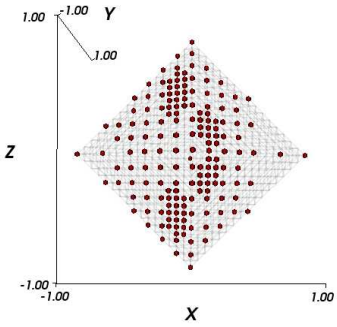

(d) wavelets: τ = 0.005 (103)

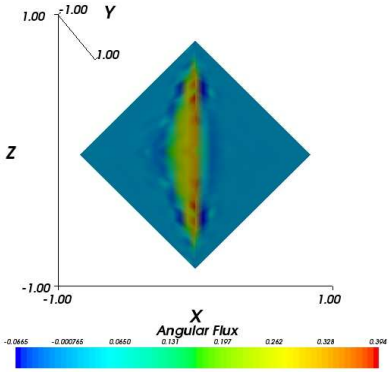

(e) Angular flux: τ = 0.001

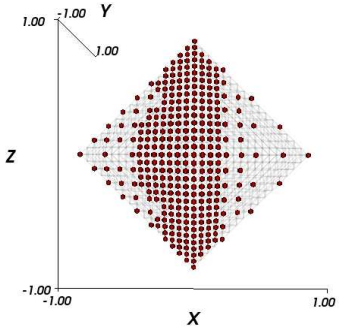

(f) wavelets: τ = 0.001 (192)

Figure 6.11: Angular flux profiles and wavelet concentration plots taken at the position (0,1.35) of the box source problem. The dots represent the collocation points associated with the wavelets included in the wavelet expansion. The number of wavelets used in the expansions are denoted by (.).

oscillations. There is also no detail in the $+x-y$ octant of the sphere. The reasons for this poor representation is shown by the wavelet position plot which reveals that no high order wavelets have been included in the expansion and that almost no resolution was applied to the $+x-y$ octant. In direct contrast to this the angular flux resulting from the tolerance $\tau = 0.001$ shows the peaked strip to be well constructed and free from large oscillations. However, the positions of the wavelets constructing this approximation reveal a heavy concentration of functions to have congregated around the $-y$ hemisphere. It is therefore possible that while the approximation is accurate, the low tolerance has enabled the smaller contributing wavelets to enter the expansion. Once again the results show the tolerance $\tau = 0.005$ to be a suitable compromise between the two extreme values. The vertical strip is again well constructed, free from large oscillations and compares closely to the flux of the $\tau = 0.001$ calculation. However in this calculation the higher order wavelets are shown to congregate only over the important areas of the sphere, namely along the vertical band of flux, thus making the calculation highly efficient.

Figure 6.12 presents a similar set of angular flux profiles obtained from the three tolerance values sampled within the problem's source region at position (1.55,1.35). Again the figure has included the position plots of the wavelets used in the three angular flux expansions. The three plots show the angular flux to be fairly uniform in magnitude about the z axis and to contain large maxima around the poles of the sphere. There is also a higher concentration of particles directed in the $+x-y$ octant of the sphere due to the slight off centre positioning of the sampling position.

Once again the profiles agree with the previous findings. The calculation using the tolerance $\tau = 0.01$ has poorly reproduced the maxima about the poles of the sphere in which large oscillations have formed. The plot corresponding to this approximation shows that only a sparse scattering of high order wavelet have formed around the poles of the sphere. This again indicates the tolerance is too high which has in turn excluded some important wavelets from the expansion. In contrast to this, the tolerance $\tau = 0.001$ has produced a well formed flux with no oscillations forming about the pole. However, as with the previous example, a large proportion of the high order wavelets have been included in the angular expansion. It is therefore highly possible that the tolerance (being too low) has again allowed to many of the non contributing wavelets to enter the angular expansion.

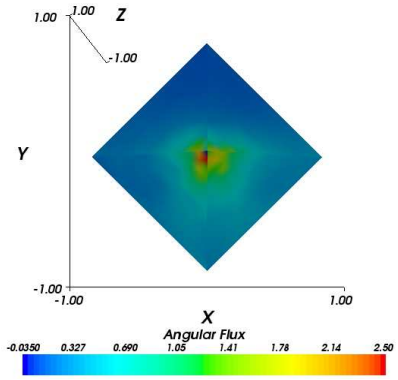

(a) Angular flux: τ = 0.01

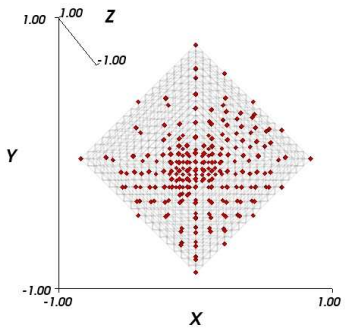

(b) wavelets: τ = 0.01 (154)

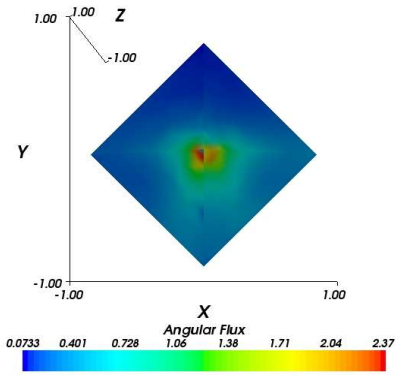

(c) Angular flux: τ = 0.005

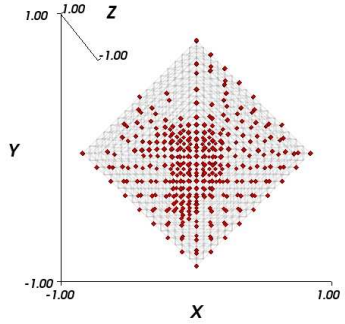

(d) wavelets: τ = 0.005 (225)

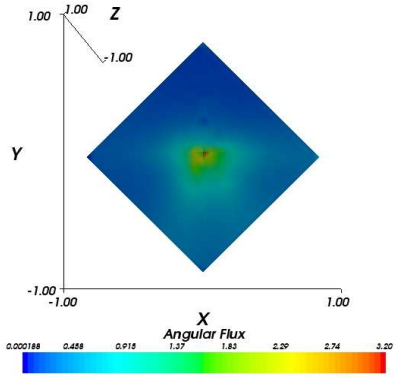

(e) Angular flux: τ = 0.001

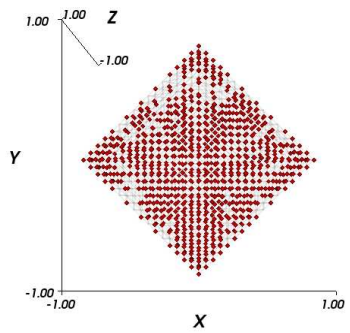

(f) wavelets: τ = 0.001 (518)

Figure 6.12: Angular flux profiles and wavelet concentration plots taken at the position (1.55,1.35) of the box source problem. The dots represent the vertexes associated with the wavelets included in the expansion. The (.) denotes the number of wavelets used in the angular expansion.

The results for the tolerance $\tau = 0.005$ have again performed well in terms of both accuracy and efficiency. The peak around the poles are well formed and contain only minor oscillations. However it is also shown that the high order wavelets have concentrated only around the poles of the sphere. Therefore the high resolution has been applied to only the regions where it is necessary.

6.3.2 Straight Duct Problem surrounded by a Purely Absorbing Region

This problem has been designed to demonstrate the ability of the adaptive wavelets to produce accurate angular flux approximations for demanding problems involving ducts containing voids. Both the octahedral and hexahedral linear wavelets are investigated and their solution's accuracy are compared to the solutions of EVENT.

The problem involves a straight duct surrounded by a pure absorbing media, see [49] page 449. The domain is presented in figure 6.13 and the corresponding source and material properties are listed in table 6.3. The problem domain is rectangular of size 36cm $\times$ 28cm. The central 6cm $\times$ 6cm square contains an isotropic source of strength 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and hosts a purely absorbing material with a cross-section $\sigma_t = 0.5$. Two near void ducts, with absorbing cross sections $\sigma_t = 0.001$, of width 6cm extend from the source to the upper and lower boundaries. The regions surrounding the ducts contain a purely absorbing material with a cross-section $\sigma_t = 0.5$. Vacuum boundaries are prescribed along all boundaries of the problem domain. A finite element mesh comprising 900 quadrilateral elements and 961 nodes was used to discretise the spatial domain. The calculations were performed with an adaptive LW_4 angular expansion. Error measure 2 with the tolerance $\tau = 0.005$ was employed to locate the areas of angles requiring high resolution.

The solution to this problem using the adaptive LW_4 angular discretisation is shown in figure 6.14. The profile displays a flux peak over the central source region reaching approximately 1.75 neutrons.$\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The strong pure absorber is shown to cause a sharp attenuation of the particles. However, the scalar flux declines at a more gradual rate

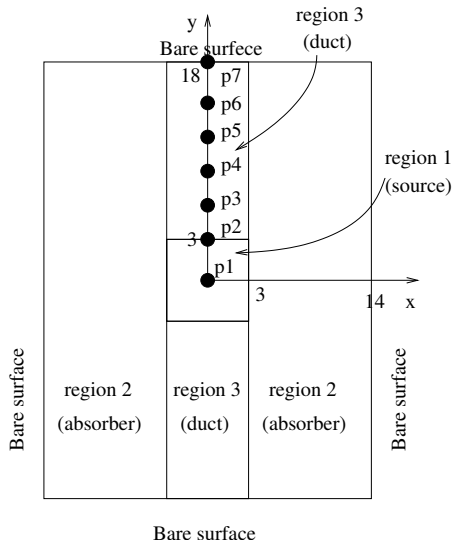

Figure 6.13: Diagram of the straight duct problem.

Region	Source $(cm^{-2}s^{-1})$	$\sigma_a\ (cm^{-1})$	$\sigma_s\ (cm^{-1})$
1	1.0	0.5	0.0
2	0.0	0.5	0.0
3	0.0	0.0	0.0

Table 6.3: Source and material properties of the Straight Duct Problem.

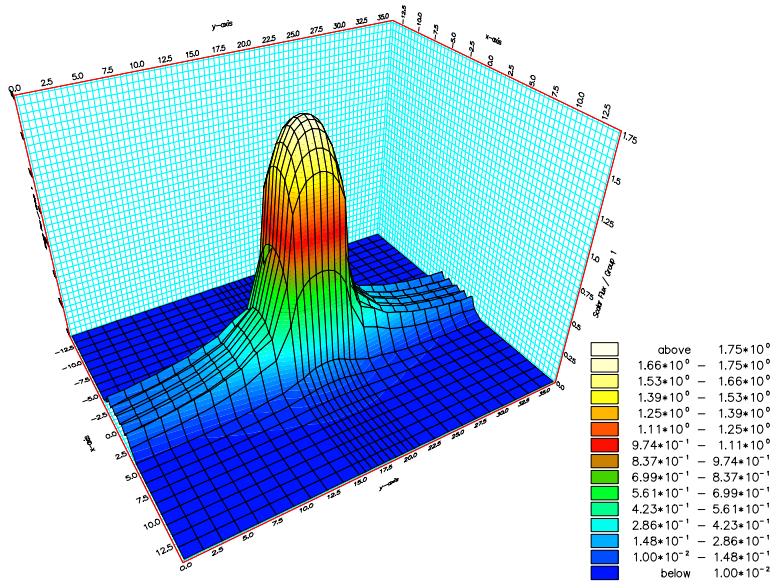

Figure 6.14: Scalar flux profile of the adaptive octahedral wavelet solution of the straight duct problem using $\tau = 0.005$.

through the ducts as particles stream through the void without collision. This solution is consistent with solution provided by even-parity P_{21} as shown by the graph in figure 6.15. The graph shows the scalar flux profile from the centre of the source, through the duct to the mid position of the top boundary. Over the source and duct regions the two solutions compare closely, however, variations do appear across the interface between the two regions. This variation is due to the large differences between the material properties of the source and duct regions which causes large gradients to occur in the scalar flux solutions. The variation is therefore a spatial discretisation issue and will not be considered any further.

Figure 6.16 presents a diagram showing the variation of angular expansion size within the spacial domain. The results show that adaptive wavelets have concentrated much effort in resolving the source and duct regions. The source region requires a large number of wavelet functions. That is because the problem's geometry generates complex flux distributions containing maxima about the poles that require high resolution to be applied to most regions of the sphere. To resolve the source region accurately approximately 200 wavelets were required at each spatial node. The result also reveals that the interface be-

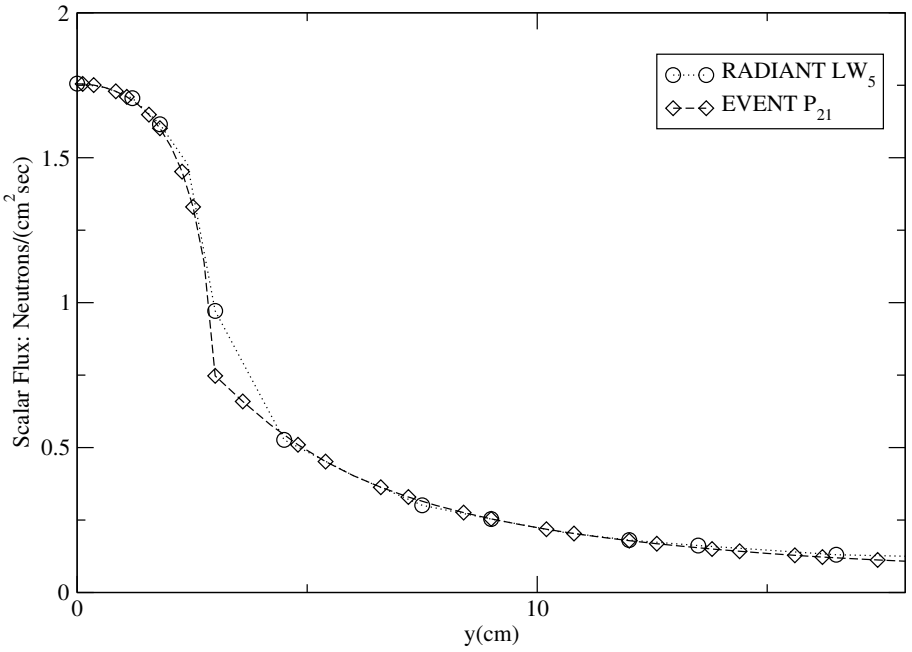

Figure 6.15: A line plot of the scalar flux solutions through the positions (0,0) to (0,18). The adaptive octahedral wavelet solutions is compared to the even-parity P_{21} solution.

tween the source and duct region were the most expensive to approximate by requiring over 300 wavelets at each finite element node. However, this figure dramatically reduced when resolving the angular flux at increasing distance from the source along the duct. This suggests that adaptivity is working efficiently as redundant wavelets have been omitted from the calculation as the flux becomes more directional. In fact, at the interface between the boundary and the duct only 170 -180 wavelet were required for an accurate representation. It is also shown that the regions deep within the absorber acquired no further resolution than that provided by the initial scaling functions. This is highly important as the scalar flux is negligible within these regions and so the discretisation has avoided wasting effort by providing it extra resolution. In total, an average of 288 wavelet functions were required to resolve this problem. The time for solving the adaptive discretised equations will therefore be between 22.1% and 47.0% of the time to solve the full non-adaptive equations.

Figure 6.17 shows seven angular flux profiles, viewed from the negative y direction, sampled at 3cm intervals through the source and duct regions as illustrated in figure 6.13. The

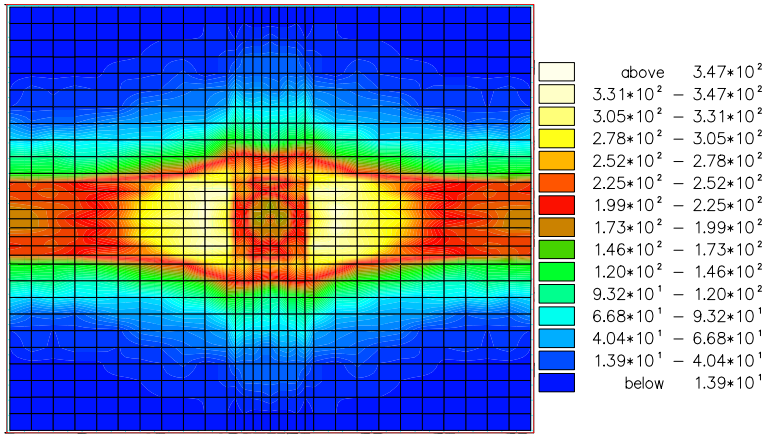

Figure 6.16: Diagram showing the adaptive octahedral wavelet expansion size over the nodes of the straight duct problem's spatial grid. τ = 0.005

figure also includes the wavelet collocation plots corresponding to the wavelets that were used in each of the angular expansions.

Figure 6.17a presents the angular flux positioned over the centre of the source region. The result shows the flux to be evenly distributed about the z axis and to contain high maxima about the poles of the sphere. The wavelet expansion corresponding to this angular flux is shown to have required the vast majority of the wavelets up to the third shell. The result also shows that the high resolution, forth shell wavelets to have been placed only in the areas where high resolution was required - namely the regions containing large gradients. In total, the number of wavelets required for this expansion was 204 - that is 33.3% of the available wavelets. Figure 6.17b shows the approximated angular flux positioned over the interface of the source and duct regions. At this position a complex angular flux is generated in the $+y$ hemisphere. The regions requiring high resolution are located close to the x axis as these angles form large gradients due to a low flux in the $-y$ hemisphere. The wavelet expansion at this position required 288 functions - that is 47.1% of the complete expansion set. The angular flux corresponding to position 3 is presented in figure 6.17c. At this position the flux becomes highly directed over the $+y$ angles due

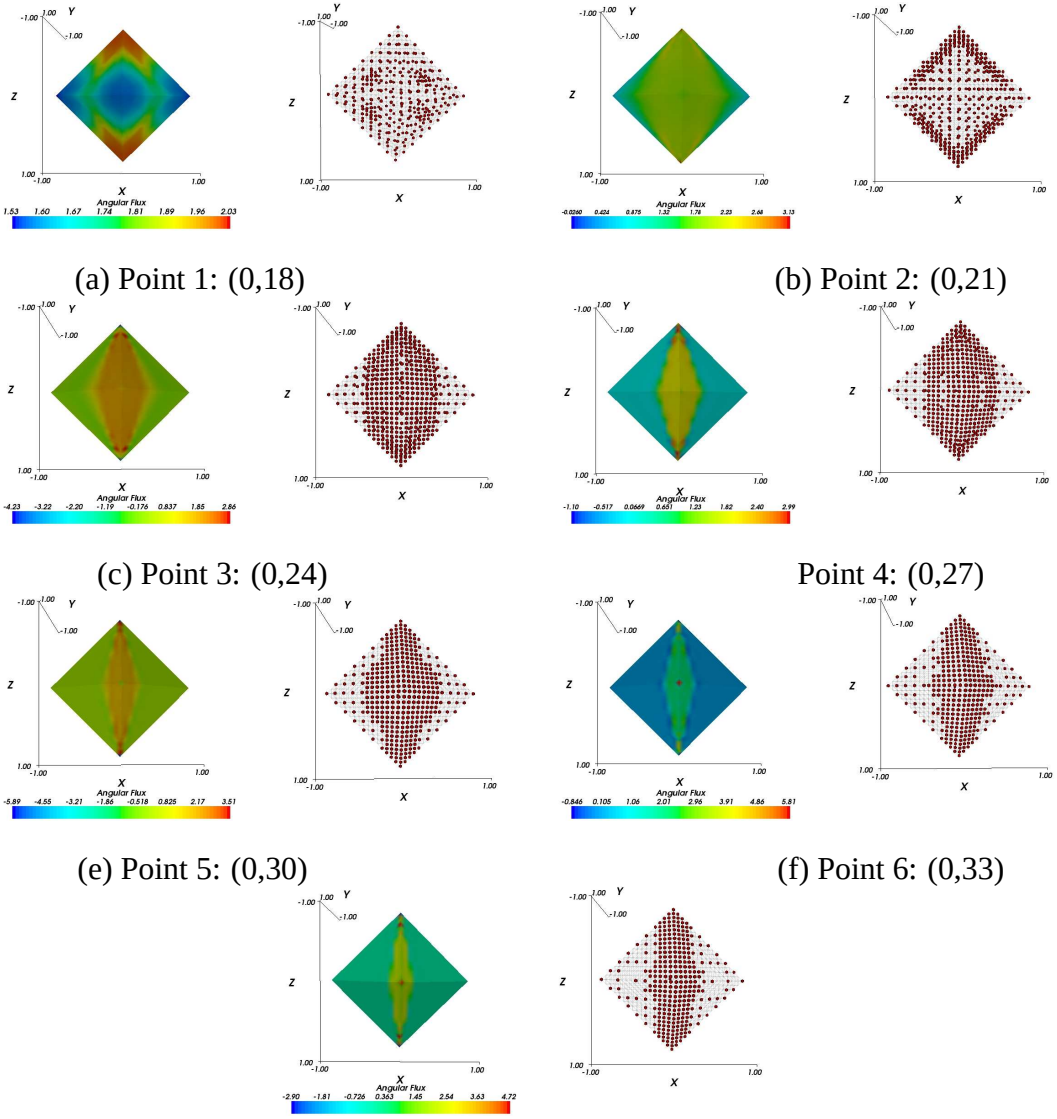

(a) Point 1: (0,18)

(b) Point 2: (0,21)

(c) Point 3: (0,24)

Point 4: (0,27)

(e) Point 5: (0,30)

(f) Point 6: (0,33)

(g) Point 7: (0,36)

Figure 6.17: Angular flux profiles and wavelet concentration plots sampled at the positions: Point 1:(0,18), Point 2:(0,21), Point 3:(0,24), Point 4:(0,27), Point 5:(0,30), Point 6:(0,33), Point 7:(0,36). The dots represent the vertexes associated with the wavelets included in the expansion.

to particles streaming through the duct. The flux's maxima has also transferred from the poles to a vertical strip along the x axis in the $+y$ direction. The results show that the adaptivity was successful in adjusting its resolution to give high order representation to these new regions where it was required. In total 322 wavelets were required to fully represent the angular flux (52.2%). The angular flux over position 4 is presented in figure 6.17d. Similar to the previous result the flux is highly directional and forms a vertical band over the x axis directing particles in the positive y direction. However, due to the position being sampled further from the source the peaked band is narrower is shape. The wavelets expansion has again successfully followed the flux's pattern by concentrating over the same (but narrower) band of angles. This in turn has and led to the number of wavelets required to represent this angular flux to decreased to 286 (46.7%). The flux profiles over the final three positions, points 5, 6 and 7, are presented in figures 6.17e, 6.17f and 6.17g respectively. The results follow a similar pattern to that of the angular flux over position 4. Each profile shows the flux to be confined to the angles forming a vertical band over the x axis, directing particle in the $+y$ direction. However, the effect of the sampled positions being at further distances from the source is that these bands become narrower in shape. This has in turn allowed the adaptivity to use less wavelets in the construction of each flux when the distance from the source is increased. To construct the angular flux at positions 5, 6 and 7, 224 (36.6%), 204 (33.3%) and 183 (29.9%) wavelets were required in the expansion, respectively.

An issue that has been brought to light by these results is the presence of Gibbs oscillations forming in the angular flux approximations. In particular, heavy oscillations are shown to manifest over the pole regions of the angular flux approximation at positions 3 and 5. In both approximations heavy oscillations have appeared due to the flux containing large gradients at the poles - in these situations Gibbs oscillations tend to be at their most damaging. However, the presence of Gibbs oscillations is not always as damaging as they may appear. For example, the area in which the flux is negative is small in comparison to surface of the sphere. Therefore the negative fluxes have only a minor influence in the scalar flux solutions.

Hexahedral wavelets

In order to demonstrate the adaptive hexahedral wavelets abilities, the method was used to approximate the angular flux for the straight duct problem. The calculation was allowed a maximum of 3 shells of wavelets, allowing up to 868 angular basis functions per spatial node. Error measure 2 with $\tau = 0.15$ was employed to locate the areas of angle requiring high resolution. Figure 6.18 presents the adaptive hexahedral wavelet's scalar flux solution together with the diagram illustrating the number of wavelets used in the angular expansion over the nodes of the spatial mesh. The scalar flux shows a similar profile to the solution obtained from the octahedral wavelets. In fact the line plot through the domain's centre shows the two wavelet solutions to be almost identical, this is presented in figure 6.19. The diagram displaying the spatial nodes' expansion size also reveals similar patterns seen with the octahedral wavelets. Again, the centre of the domain required the highest resolution in which the full 868 wavelet functions were used in the approximation. Large wavelets expansions also concentrated along the ducts. However, as previously seen with the octahedral wavelets, the expansion size along the ducts decreased as the distance from the source increased. The adaptive hexahedral wavelets were also successful in avoiding the placement of unnecessary resolution over the absorber regions that acquired negligible fluxes. Over all the average number of wavelets per spatial node was 202, that is 23.3% of the available wavelets.

Figure 6.20 presents the angular flux profiles at positions 4 and 7 (as shown in figure 6.13). The diagrams clearly show the vertical peaked band of flux resulting from particles streaming through the ducts. Once again the fluxes show a narrowing of the band as the distance from the source region is increased. The wavelet expansions have successfully adapted in order to apply an appropriate resolution to all regions of the sphere. The results show that no wavelets directing flux in the $-y$ direction have been included in the expansions. It is also shown that the angular flux closest to the source region required the majority of the high ordered wavelets centred about the x axis. However, for resolving the flux further from the source, the contracting peak has allowed many of these wavelets to be omitted. The results also show the hexahedral wavelet solutions to have developed Gibbs oscillations. Not surprisingly the oscillations manifest over the same areas of angle

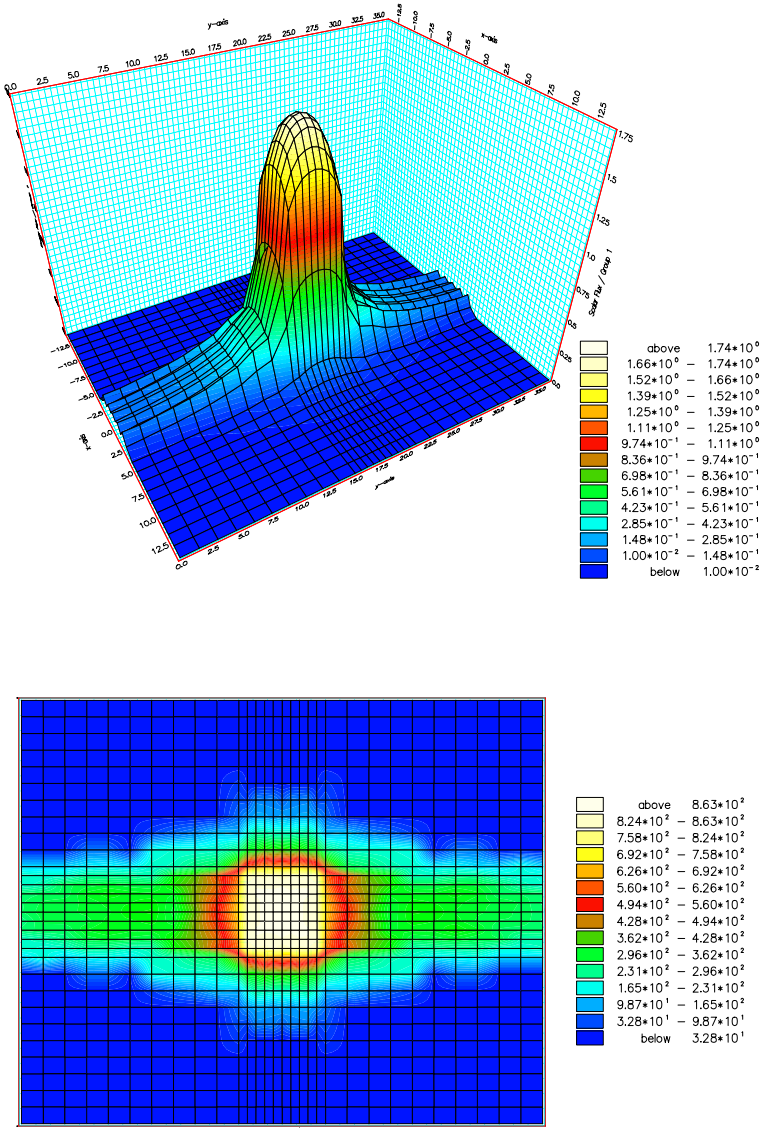

Figure 6.18: Top: Scalar flux profile of the adaptive hexahedral wavelet solution using τ = 0.15. Bottom: Plot displaying the number of hexahedral wavelets used over the spatial mesh

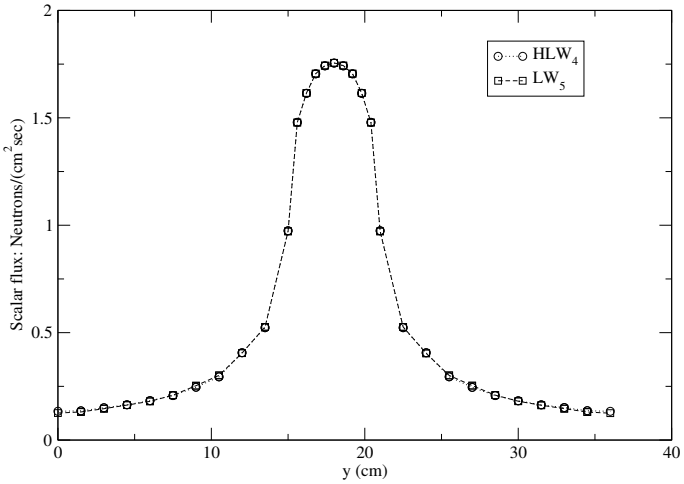

Figure 6.19: Scalar flux profile of the linear octahedral and hexahedral wavelet solutions measured along the line between the coordinates (0,0) and (0,36).

as observed with the octahedral wavelets. However the oscillations do appear to be less severe in the hexahedral wavelet solution. The number of hexahedral wavelets employed to resolve positions 4 and 7 were 364 (41.9% of the possible wavelets) and 288 (33.1% of the possible wavelets), respectively.

6.3.3 The dogleg duct problem

In order to establish the strengths and abilities of the adaptive wavelet discretisation, a problem where neutron streaming occurs in two directions along a near void duct is considered. For this demonstration the linear octahedral wavelets are used. Various tolerance setting are considered in order to show the value's influence on the adaptive solutions. Solutions for comparing the wavelets accuracy are provided by EVENT and the discontinuous Galerkin - Discrete ordinate DG-DO codes.

The problem is generally known as the dog-legged duct problem and is referenced in page 465 of [49]. The problem domain is presented in figure 6.21 and the corresponding source and material information is listed in table 6.4. The problem's upper quarter (note it has symmetry along both the x and y axes) has a rectangular domain of size

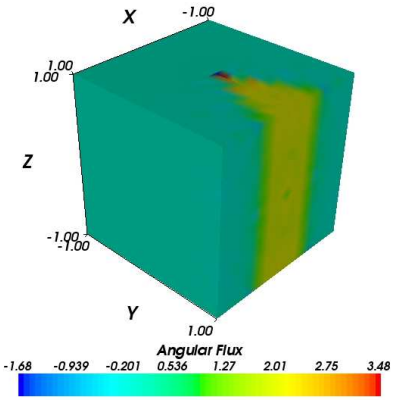

(a) Point 4: (0,27)

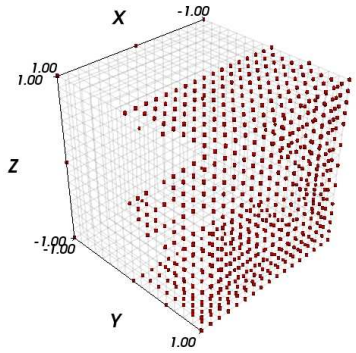

(b) Point 4: (0,27)

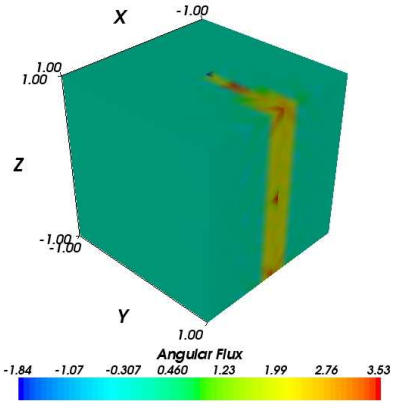

(a) Point 7: (0,36)

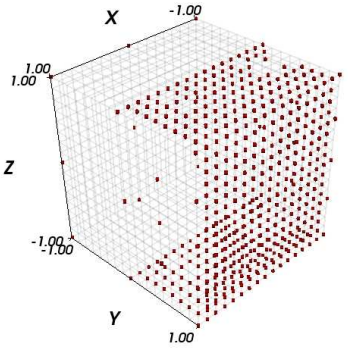

(b) Point 7: (0,36)

Figure 6.20: Angular flux profiles and wavelet concentration plots from the hexahedral wavelet calculations. Points 4 (top) and 7 (bottom) along the ducts were measured.

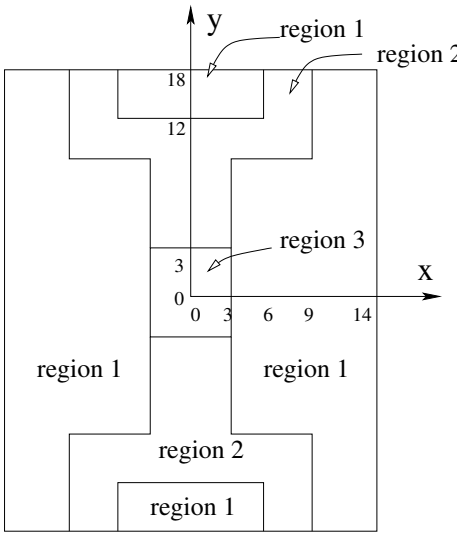

Figure 6.21: Diagram of the dog-leg problem.

14×18cm. The bottom left square region contains a 3×3cm source of intensity 1.0 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. A duct of width 3cm runs from the source for 9cm up along side the y axis. The duct then takes a 90 degree turn to the right and proceeds a further 6cm. The duct then takes one final 90 degree turn to the left and continues to the problem's top boundary. An absorbing material surrounds the duct with a cross-section $\sigma_t = 0.5\mathrm{cm}^{-1}$. Bare surface boundary conditions are applied to the top and right boundaries. A mesh consisting of 1,296 nodes and 1,225 quadrilateral elements was used in order to discretise the spatial domain. Calculations were performed with the adaptive linear octahedral wavelets enabling a maximum expansion of 612 angular basis functions. Error measure 2 was employed with the tolerance set to $\tau = 0.025$. In order to help clarify the solution's accuracy a number of even-parity calculations were also performed. In these calculations the spatial and angular resolution was increased to very high levels in an attempt to obtain a converged solution.

Figure 6.22 presents the scalar flux solutions of the adaptive wavelets and even-parity. The even-parity calculation used a P_{19} angular expansion and a very fine FE spatial mesh consisting of 122,500 elements. The two solutions appear to be in close agreement. The scalar flux peaks over the centre of the source region at approximately 4.9 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Both results show a high attenuation of the particles through the absorbing

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	0.0	0.5	0.0
2	0.0	0.005	0.0
3	1.0	0.005	0.0

Table 6.4: Source and material properties of the dog-leg problem.

media. Through the first section of the duct the flux declines at a slower rate as particles stream through the void without interacting with the material. Upon reaching the first turn, the particles continue into the absorber and are quickly absorbed. The results also show a small number of particles avoid collision with the absorber and reach the domain's boundary via the 3 segments of the duct.

Figures 6.23 and 6.25 present scalar flux profiles through the lines (0,0)-(0,18) and (0,18)-(14,18) of the scalar flux solutions obtained from three even-parity and wavelet calculations. The even-parity calculations used P_{15} and a 30,625 FE spatial mesh, a P_{19} and a 30,625 FE spatial mesh and P_{19} and a 122,500 FE spatial mesh. Figure 6.23 plotting the scalar flux from the domain's centre to the top boundary shows the wavelet solution peaks at approximately 4.95 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ - which is some way above the peak of 4.3 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ obtained by the even-parity P_{15} calculation. However, increasing the even-parity's angular approximation to P_{19} returns a scalar flux that peaks at 4.6 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This increase in angular resolution has resulted in halving the original variation - this is illustrate by the corresponding error plots presented figure 6.24. This gives a strong evidence that the even-parity calculation has failed to converge in angle, even with this high level of P_N. The graph also demonstrates that an additional increase in the spatial resolution also produces an increase in scalar flux peak. The Even-parity P_{19} using the spatial mesh consisting of 122,500 elements has produced an increased scalar flux peak of 4.75 particles $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Again this implies that even-parity has also failed to converge in its spatial approximation.

This pattern of slow convergence is also shown in the results of figure 6.25. The successive increases in the even-parity's angular and spatial resolution produces lower scalar fluxes measured at the boundary containing the duct's mouth. This is highlighted by the

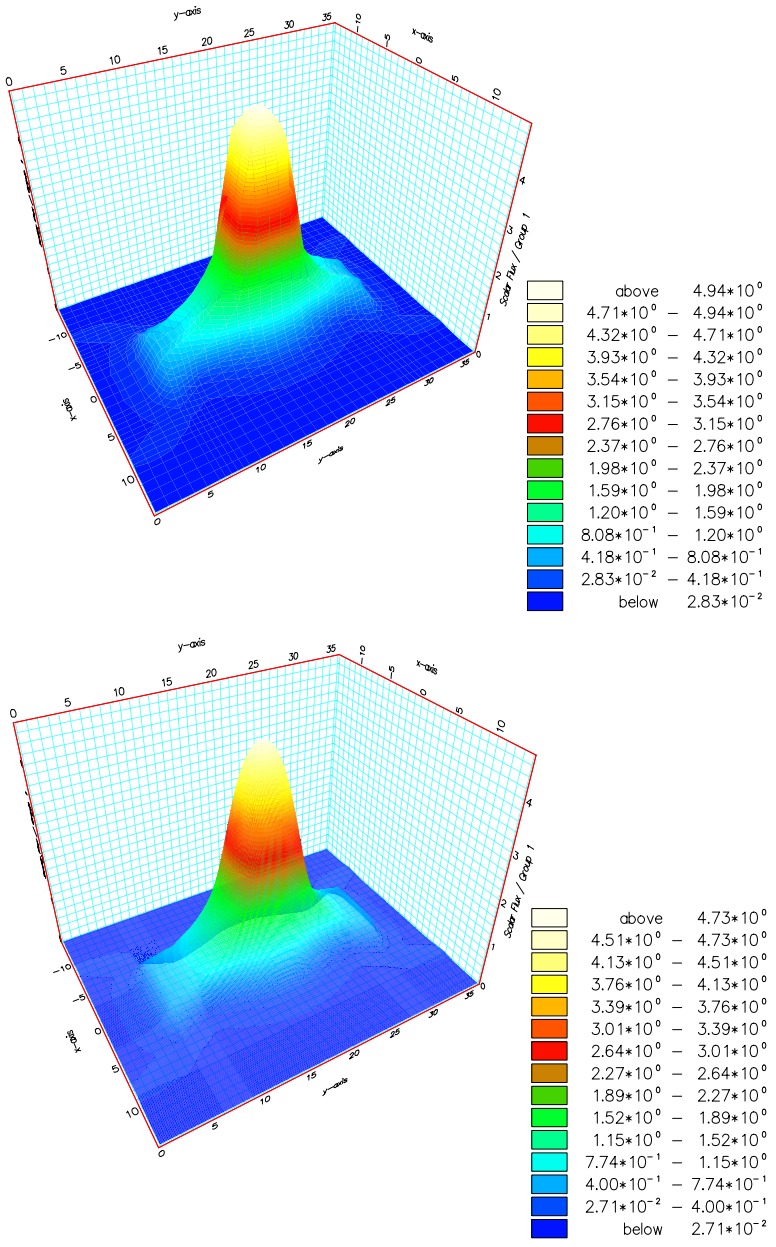

Figure 6.22: Scalar flux profile of the adaptive solution using $\tau = 0.025$ (left) and EVENT P_{19} on a high resolution spatial mesh (right).

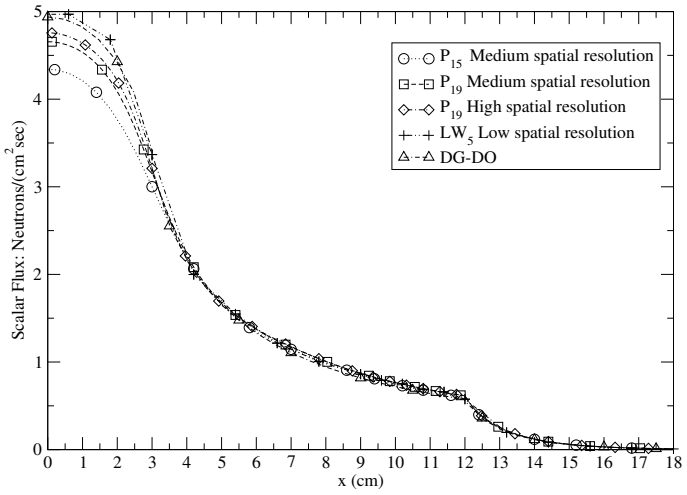

Figure 6.23: A line plot of the scalar flux solutions through the positions (0,0) to (0,18) of the dogleg problem. The linear octahedral wavelet solutions are compared to the even-parity P_{19} solution.

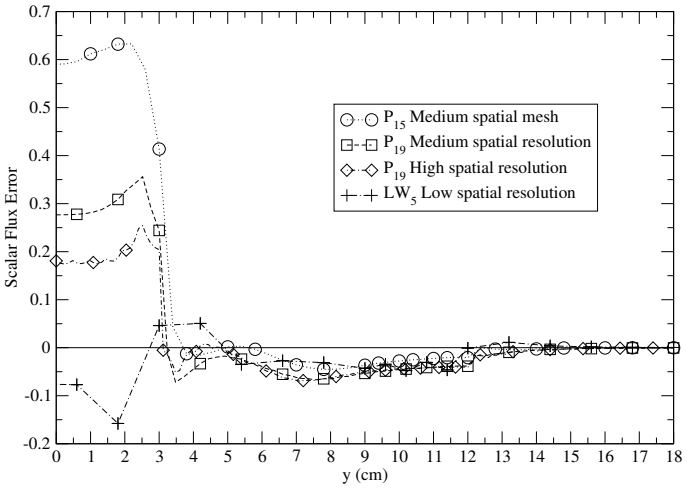

Figure 6.24: A line plot of the scalar flux errors through the positions (0,0) to (0,18) of the dogleg problem. The linear octahedral wavelet errors are compared to the even-parity P_N solutions.

corresponding error plots of in figure 6.26. These two graphs show evidence that the even-parity is highly dissipative for this problem.

Both graphs of figures 6.23 and 6.25 reveal that increases in even-parity's angle and spatial

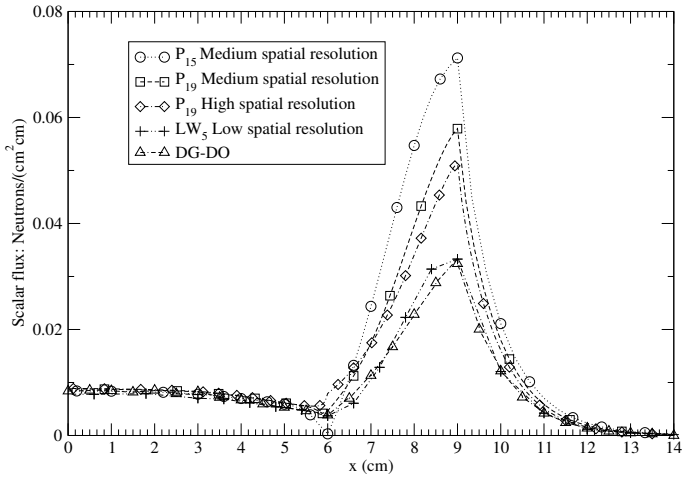

Figure 6.25: A line plot of the scalar flux solutions through the positions (0,18) to (14,18) of the dogleg problem. The linear octahedral wavelet solutions are compared to the even-parity P_{19} solution.

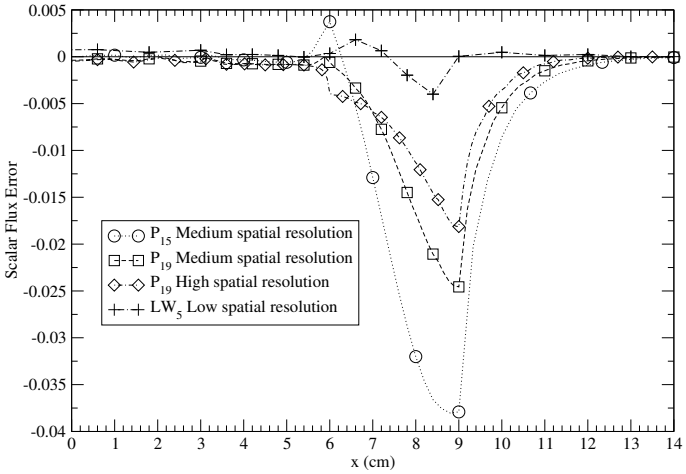

Figure 6.26: A line plot of the scalar flux errors through the positions (0,18) to (14,18) of the dogleg problem. The linear octahedral wavelet errors are compared to the even-parity P_N solutions.

resolution produces solutions converging toward the wavelets calculation. However, in order to establish a benchmark solution a discrete ordinate, discontinuous Galerkin finite element (DG-DO) calculation was performed using a 10,000 ordinate S_N quadrature set over a spatial mesh consisting of 14,700 DG elements. The scalar flux solutions have been included in the graphs presented in figures 6.23 and 6.25. The results reveal close profiles between the wavelet and DG-DO solutions, confirming that the wavelet solution is in close agreement with the benchmark DG-DO solution. These results also confirm the even-parity solutions to be highly dissipative, underestimating the true solution over the source region and overestimating the flux at the boundaries.

Figure 6.27 shows the spatial variation of the number of wavelet basis functions used. This figure clearly shows the source region to be the most difficult to resolve. It requires the full set of 612 wavelet functions. The figure also shows the large number of wavelets needed to resolve the duct regions. Note here that even the bends in the duct are clearly visible. Again, as with the previous numerical example, the absorbing regions with negligible flux have the least angular resolution. In total the number of angular basis functions used in this calculation was 11,572, that is 11.1% of the total angular basis functions available. The time for this calculation can therefore be expected to be between 1.2% and 11.1% of the time to resolve the full non-adaptive equations.

Figure 6.28 presents six angular flux approximations sampled at various positions along the problem's duct (as shown in figure 6.21). The diagrams also indicate the density of the wavelets on the sphere. For the first four plots large Gibbs oscillations were found to dominate the angular flux scale. These plots have therefore had their scale restricted in order that the detail of the angular flux may be seen.

Figures 6.28a and 6.28b present the angular fluxes at positions 1 and 2, respectively. These are positioned in the first section of the duct, before the bend, and have a direct line of sight with the source. The angular flux profiles show highly directed fluxes travelling in the negative y direction. The two plots reveal a narrowing of the peak in angular flux as the distance from the source is increased. The diagrams demonstrate the wavelets to have adjusted to this by their exclusion of high resolution in the $-x-y$ regions on the sphere. This in turn has reduced the expansion from 254 wavelets (41.5% for the

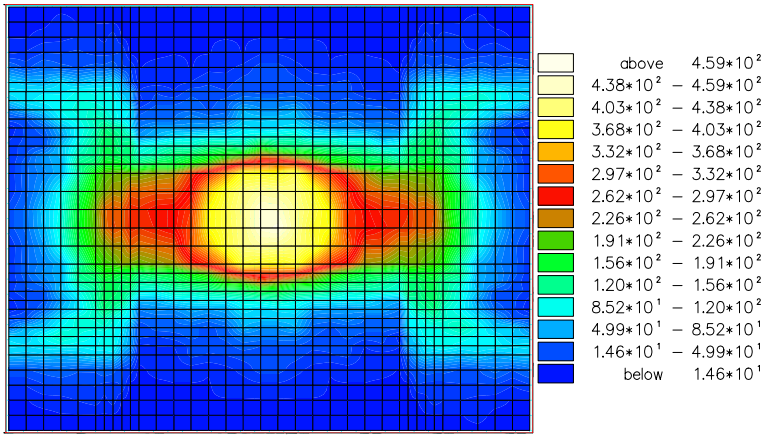

Figure 6.27: Diagram plotting the number of octahedral wavelet basis functions used in the angular expansions over the nodes of the dogleg duct problem's spatial grid. τ = 0.025

maximum available functions) to resolve the flux at position 1 to 236 (38.5%) wavelets to approximate position 2.

The angular fluxes at positions 3 and 4 are presented in figure 6.28c and 6.28d, respectively. Positioned within the mid section of the duct and out of direct line with the source the two fluxes display highly directional profiles in the $+x - y$ octant of the sphere. The flux at position 3 has a well formed peak that has required the majority of the high order wavelets positioned over this range of angles (142 wavelets - 23.3%). The profile of the angular flux at position 4 reveals a reduced peak which results from the particles having to travel further through the absorber region in order to reach this point. The wavelet have performed well by adding high resolution to only the directions close to the $+y$ axis, therefore requiring just 82 (13.2%) of the available wavelet functions to resolve the angular flux distribution.

Figure 6.28e and 6.28f present the two fluxes at positions 5 and 6, respectively. Both flux distributions display highly directional profiles. This has allowed the vast majority of the wavelet to be excluded from the flux approximation, requiring only 98 (16.0%) wavelets

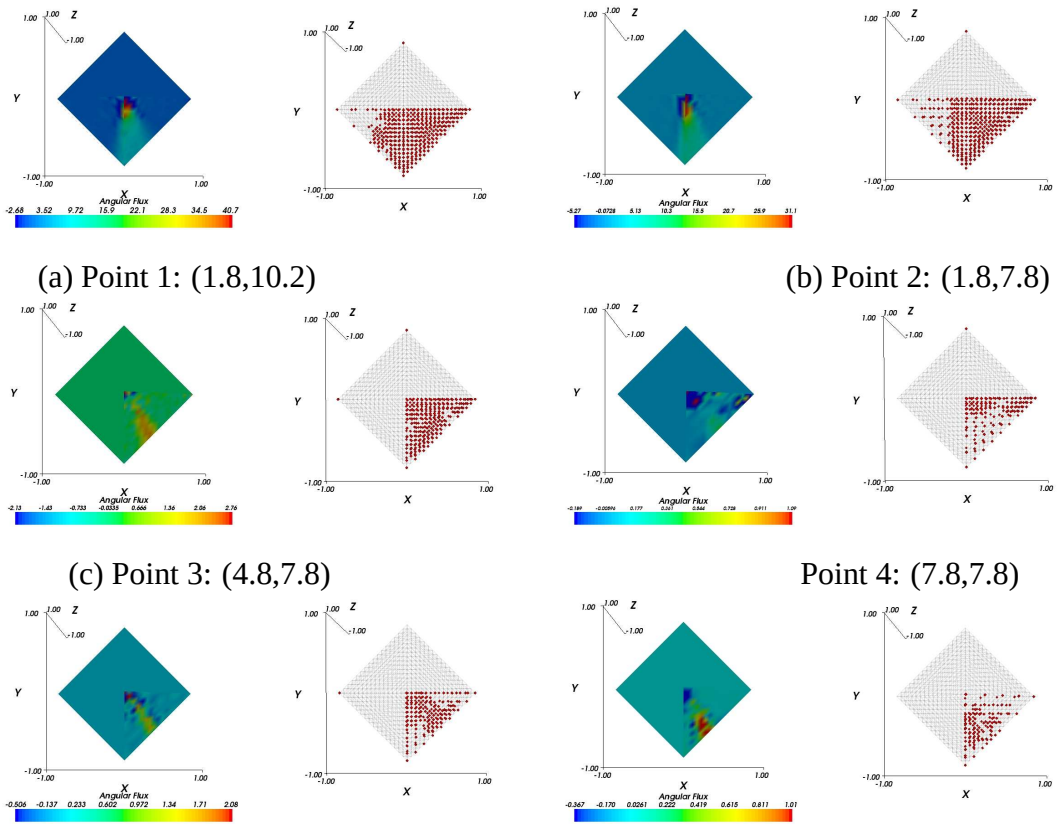

(a) Point 1: (1.8,10.2)

(b) Point 2: (1.8,7.8)

(c) Point 3: (4.8,7.8)

Point 4: (7.8,7.8)

(e) Point 5: (7.8,4.8)

(f) Point 6: (7.8,1.2)

Figure 6.28: Angular flux profiles and wavelet expansion positions sampled at the positions: Point 1:(1.8,10.2), Point 2:(1.8,7.8), Point 3:(4.8,7.8), Point 4:(7.8,7.8), Point 5:(7.8,4.8), Point 6:(7.8,1.2). The dots represent the vertexes associated with the wavelets included in the expansion. $\tau = 0.025$

for position 5 and 68 (11.1%) wavelets for position 6.

In order to investigate the solution's sensitivity to the error measure's tolerance the angular flux was recorded at position 5 using various adaptivity tolerances τ. Figure 6.29 presents the angular flux and wavelet profiles viewed from the negative y direction obtained using $\tau = 0.005, 0.01, 0.025$ and 0.05. The lowest tolerance $\tau = 0.005$ produces the well defined peak in angular flux which accurately depicts the particle distribution emanating directly from the source region. However, in achieving this highly accurate approximation the complete set of wavelets directed over the $+x - y$ octant were included in the expansion. A similar situation arises with the tolerance level $\tau = 0.01$. The vertical strip of flux is clearly visible but again the wavelet expansion included all but a few of the higher order functions in the $+x - y$ octant of the sphere. The tolerance value of $\tau = 0.05$ has managed to significantly reduce the number of wavelets. The higher order wavelets that are included in the expansion are concentrated along the peak of the angular flux. However, some of the wavelets along this peak are excluded causing a poor representation of the flux around this region. The tolerance $\tau = 0.025$ however appears to be the best of the sampled values. As already demonstrated the high resolution is applied only to vicinity of the angular flux maximum.

This problem has proved to be particularly demanding. This was evident from the even-parity solutions failing to converge in space and angle, even when high resolutions were used. The wavelet solution appears to be close to the benchmark solution, provided by the DG-DO calculation, which has again showed the wavelet's ability to adapt to the problem by applying high resolution to the areas of angle that are important. This demonstration also highlights the problem associated with selecting an appropriate tolerance for the error measure.

6.4 Conclusions

In this chapter a new method for self adaptive angular discretisation of the first order Boltzmann transport equation has been presented. The method employs hierarchical an-

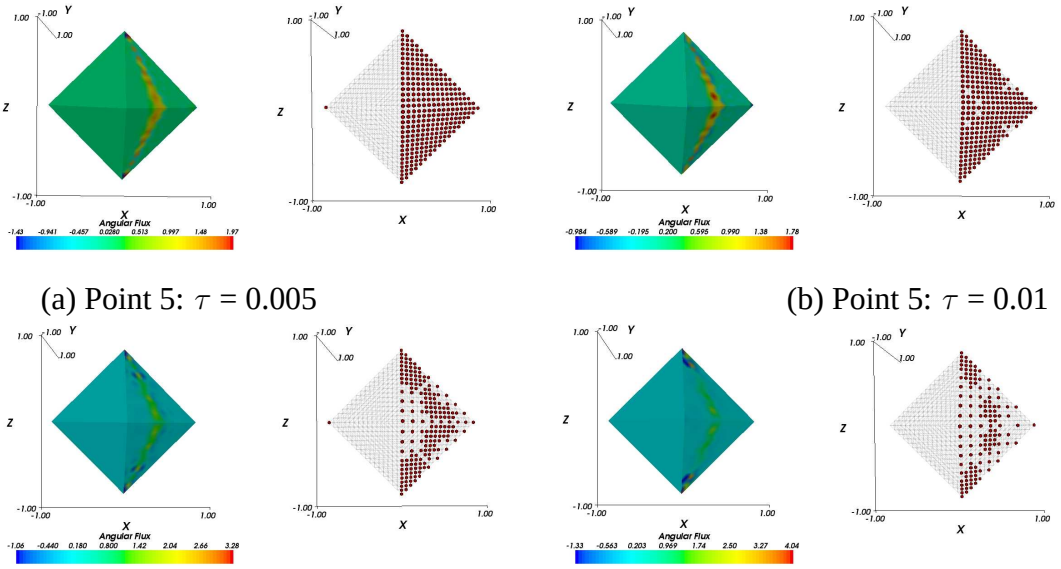

(a) Point 5: τ = 0.005 (b) Point 5: τ = 0.01

(c) Point 5: τ = 0.025 Point 5: τ = 0.05

Figure 6.29: Angular flux profiles and wavelet concentration plots sampled at point 5 (7.8,4.8). (a) τ = 0.005, (b) τ =0.01, (c) τ = 0.025 and (d) τ = 0.05

gular expansions using second generation spherical wavelets to represent the directional variable of the angular flux. The procedure used the hierarchy and compact support properties of the wavelet functions to locate areas on the unit sphere in which the flux is under resolved and to add functions that increased resolution to these regions. The methods described have been applied to a number of two dimensional, steady state, mono energetic neutron transport problems. These problems were designed to generate complex angular flux distributions in order to test the methods capabilities. The method's results were compared to the results from two established radiation transport methods. One method solved the even-parity form of the Boltzmann transport equation using spherical harmonics to represent the angular variable and finite elements in space. The second method uses an S_N representation of the angular dependence of the first order form of the transport equation and a discontinuous Galerkin spatial discretisation. The demanding numerical examples demonstrated that the adaptive wavelets are capable of resolving the angular flux accurately and efficiently. This was illustrated by the problems involving neutron streaming through voids in which the angular flux became concentrated to small regions on the sphere. The wavelets adapted successfully by applying high resolutions to only

these regions in angle. This in turn led to a significantly reduced wavelet expansion for approximating the angular flux - which in turn increased the computation efficiency for solving the problems by up to 90% (in some cases the time efficieny was more than this). The demonstration also highlighted the method to be numerically robust. This was demonstrated in the dog-leg duct example in which a weakness was exposed in the even-parity code where the solutions failed to converge in angle, even when using high order spherical harmonic expansions. The wavelets expansion did not experience this problem and went on to provided highly accurate approximation that compared closely to the discrete ordinate solution.

Chapter 7

A HIERARCHICAL PRECONDITIONED FGMRES-BASED SOLUTION METHOD FOR THE BOLTZMANN TRANSPORT EQUATION

Contents

This chapter presents a new angular hierarchical preconditioned variant of FGMRES that is specifically designed to solve the linear system generated by the discretised Boltzmann transport equation. The preconditioner is based on algebraic and geometric multigrid ideas for which the shells of any arbitrary hierarchical angular discretisation are used to generate the multigrid levels. The new solver, IVOR, based on a variant FGMRES algorithm, is also developed. This new solver offers an alternative approach to the restart method by it discarding and replacing specific vectors within the solver when the direction search space is full. The solver also allows the algorithm to partially restart the search direction space by using the latest residual to calculate new search directions. The steady state Boltzmann transport equation is discretised in angle and space using an arbitrary angular approximation and finite elements, respectively. A demonstration of the hierarchical preconditioner and new solver are provided by the solving of two neutron transport problems. A number multigrid cycles are examined for their efficiency in term of reducing the number of solver iterations required for its approximation to converge. The preconditioner is also compared to other standard preconditioning techniques. The numerical examples also examine the new preconditioner in order to investigate the solution's convergence rates (number of iterations) with respect to the number of angular multigrid levels. A demonstration of the intelligent vector replacement method is also presented.

7.1 Introduction

In this chapter a new matrix equation solver method is developed that is specifically designed for solving the system of equations formed through discretising the dimensions of the Boltzmann transport equation. Numerically approximating deterministic solutions of the Boltzmann transport equation requires the 7 dimensional phase-space comprising of space, angle, energy and time to be discretised. However, when discretising all 7 dimensions the resulting linear system can often become extremely large and badly conditioned. Therefore, in order to solve these types of systems efficiently, careful design of both approximation method and solver is required in order that robustness and fast convergence is ensured. For this purpose this chapter develops the new solver IVOR (Intelligent Vector Orthonormal Replacement). IVOR is an enhanced FGMRES algorithm which incorporates the new methods of partial restarts and search vector replacements. In addition to this, this chapter also develops a new hierarchical angular multigrid preconditioner in order to increase the IVOR algorithm's efficiency.

The IVOR algorithm, based on FGMRES [154], originates from the family of Krylov subspace solvers. This family includes the more familiar names of Conjugate gradient (CG) [155, 156] and GMRES [154]. The conjugate gradient method was developed in the 1950s independently by Hestenes, Stirfel [157] and Lanczos [158] for solving symmetric positive definite linear systems. Its establishment re-directed the development of iterative schemes (which had mainly been restricted to the Jacobi and Gauss type solvers) and began the new generation of Krylov subspace methods. However, CG was not immediately viewed as an iterative solver. Instead the method of conjugate gradients was considered to be a direct solver as the exact solution was guaranteed in at most n iterations - n denoting the size of the linear system. The properties of conjugate gradient's use as iterative solvers were however quickly recognised by Hestenes. It was seen that almost all of the method's iterations served to reduce the residual norm and so the method could provide close approximations with a substantially reduced number of steps. Over a period of 25 years CG established itself to become a first choice iterative solver of symmetric positive definite linear systems [159]. Its popularity was particularly established in its use in solving large sparse linear systems formed through the discretisation of certain types

of partial differential equations (PDEs).

The establishment of CG was soon followed by alternative Krylov methods for solving unsymmetric linear systems. Some early examples include ORTHODIR [154], ORTHOMIN [154], FOM [154] and BI-CG [160]. However, these methods suffered from being numerically unstable and were often prone to break down. The breakthrough in the Krylov methods came with the invention of the GMRES algorithm, introduced by Saad and Schultz in 1986 [161], which was designed to avoid such weaknesses. The algorithm is a generalisation of MINRES [162] which in turn was a generalisation of CG for solving symmetric indefinite linear systems. The advantage of GMRES is that it is guaranteed to find the approximation with the minimum residual norm from the space of all search directions. It is also numerically robust. However, the algorithm includes an orthogonalisation of the search directions, using the Arnoldi procedure [163], and requires all search directions to be stored in memory. High computational costs are also incurred in orthogonalising large sets of basis search vectors. In fact, these computational costs increase linearly with the iterations of the algorithm. Techniques for reducing high computational and memory costs are to restart GMRES or use incomplete orthogonalisation [164]. Restarting is commonly employed, however, in doing so the user loses the guarantee of convergence in at most n iterations and may also see a slow down in convergence. IVOR is designed to provide an alternative option to restart when the search direction space is full. This new solver allows a vector of a full search space to be discarded and replaced with a new search direction. This new vector is then orthonormalised with respect to the remaining search directions. The advantage of this method is that specific directions may be discarded that allow better search directions to be found. This technique also retains a large proportion of the information of the previous search directions - all this information is lost with a complete restart.

Preconditioners are commonly employed within GMRES in order to maintain the efficiency of the algorithm by ensuring that convergence is found in the solution's approximation using only a small number of iterations. The preconditioned GMRES [154] contains a fixed preconditioner within the algorithm that acts as an approximation to the linear system, but which is significantly easier to solve. Another variation of the algorithm allows the preconditioner to vary at every iteration step, FGMRES [154]. This in turn gives the

algorithm total flexibility in choosing its search directions but carries the penalty that the overheads for storing the search vectors in memory are doubled. Common GMRES preconditioners employ a small number of sweeps using standard Jacobi, Gauss Seidel, SOR or SSOR stationary iterative solvers [152]. However, recently there has been considerable research made in developing alternative preconditioners with a large amount directed into the use of multigrid techniques (MG). Here, in this section, we continue this area of development. IVOR, which includes the flexibilities of FGMRES, incorporates a multigrid preconditioner based in angle. A significant advantage in using this preconditioner is that it may vary during the evolution of the solver. The solver may therefore concentrate effort in resolving specific angular moments that are important or under-resolved.

Like stationary iteration techniques, multigrid methods are developed as standalone solvers for linear systems [165, 166, 167, 156]. Over the past couple of decades MG methods have seen their use spread to numerous applications. In particular, they have generated a considerable amount of interest in solving discretised PDEs where problems of ever increasing complexity have allowed MG's unique properties to stand out above other solvers. MGs have been applied in the areas of solving convection diffusion problems [168], elliptic equations [169], computational fluid dynamics (CFD) using compressible and incompressible [170, 171] Navier-Stokes equations and nuclear reactor transport problems [172]. Part of MG's success is that the method's approximation converges quickly over both the long and short wave information of a PDE's solution [165]. This has immediate benefits over traditional stationary solvers from which fast reductions in only the short range errors are obtained. MG methods are partitioned into two categories: algebraic multigrid (AMG) [173] and geometric multigrid [174]. Geometric multigrid forms the more traditional MG approach for solving discretised PDEs. The geometric approach constructs a succession of coarsened grids using the approximation method employed to discretise an equation. Linear systems are then generated for each grid and used to reduce the error terms of an approximation in relation to its corresponding grid. Typically, short ranged errors over the fine grids are relaxed using stationary iteration solvers and long ranged errors are mitigated through solving the coarse (and small) linear system directly. Geometric MG works naturally within discretisations using regular grids and hierarchical techniques. However, AMG has recently become more useful in unstructured mesh

problems since the method is not reliant on a grid coarsening technique [173].

Multigrid's powerful ability to solve PDEs efficiently have been enhanced by their use as preconditioners for Krylov subspace solvers. In [175] MG was shown to be a powerful preconditioner for GMRES when used to solve the one and two dimensional Helmholtz equations. The results showed that as a preconditioner, multigrid calculations converged on the solution using a number of iterations that was independent of h (the element size of the regular mesh). That is, the number of iterations required to solve the equations did not increase when finer discretisation meshes were used. This property was particularly attractive since the MG method used as a standalone solver was shown to diverged in its number of solver iterations when used on finer meshes. This undesirable characteristic was therefore avoided with the use of MG as a preconditioner. Here it was also shown that the number of iterations used by the MG preconditioned Krylov solver converged with respect to the number of MG grids. A similar conclusion was found in [176] where MG was compared as both a solver and preconditioner (for GMRES and BiCGSTAB [154]) for solving singular perturbed problems. Again there were examples showing MG to be unsuccessful as a solver but successful as a preconditioner. Other examples of multigrid preconditioners for Krylov solvers include applications in the solutions of convection diffusion equations [177, 178], incompressible Navier-Stokes [170, 171] equations and radiation transport problems [179, 172, 20].

This chapter develops an angular multigrid preconditioner for the IVOR algorithm which is applied to solve the first order Boltzmann transport equation. The MG method uses arbitrary hierarchical angular approximations of the angular flux - for example spherical harmonics [15] or spherical wavelets [1, 2]. The multigrid meshes are generated by the hierarchical levels of the angular flux approximation with the remaining space, energy and time discretisations kept constant (wavelet discretisations use finite element type meshes to generate the MG levels but for P_N, where no meshes are generated, shells of spherical harmonics for each level of P_N are used). The potential benefits of this is that the preconditioner is capable of concentrating effort on specific angular moments that are considered to require the highest attention. It is also expected that the solver will posses attributes common to other MG preconditioned solvers. One such important attribute is that convergence in the number of solver iterations should be found with respect to the

number of multigrid levels.

The following sections are set out as follows. Section 7.2.1 reviews the GMRES algorithm together with its variants preconditioned GMRES and FGMRES. This section then develops new variants for FGMRES using intelligent direction updates - IVOR. Section 7.2.2 describes the multigrid methods and their application as preconditioners for the GMRES and FGMRES algorithms. Section 7.2.3 introduces the angularly discretised Boltzmann transport equation for which the angular discretisation is assumed to be hierarchical. A Streamline Upwind Petrov Galerkin (SUPG) finite element (FE) discretisation of the spatial variables is presented. This treatment of the spatial variables is consistent with the radiation transport code RADIANT [83, 84]. The angular multigrid preconditioner is described and an account into maximising efficiency of the solver's operations (in relation to solving the BTE) is addressed. In section 7.3 some numerical examples are presented. All calculations were performed using RADIANT. The aims are to illustrate the capabilities of the angular multigrid preconditioner employed within IVOR. Comparisons of convergence are made with standard stationary preconditioners using two vastly different radiation transport problems: one relatively easy and one relatively difficult. The results try to establish the effects of various multigrid cycles on the convergence rate of the solver. In sections 7.4 the chapter finishes with a conclusion.

7.2 Angular Hierarchical Preconditioned FGMRES-Based Solution Method

7.2.1 Krylov subspace solvers

GMRES was developed to approximate solutions of large, sparse, non-symmetric linear systems of the form,

$$\mathrm{A}x = b. \tag{7.1}$$

The algorithm finds an approximation to the system by continuously updating an initial

guess by transversing along search directions and finding nearer estimates to the exact solution. At each stage i of the iterative process, an approximation x_i to the system's solution is found by the projection onto a subspace $x_0 + K_i$. The vector x_0 denotes the starting point or initial guess of the exact solution, x, and K_i represents the spanned basis of search directions which is defined as the Krylov subspace for a vector v and matrix A.

$$K_i(\mathrm{A}, v) = \text{span}\{v, \mathrm{A}v, \mathrm{A}^2v, \ldots, \mathrm{A}^{i-1}v\} \tag{7.2}$$

By construction, GMRES is guaranteed to find the approximation x_i with the least square error from all the available vectors of the subspace K_i. That is, GMRES finds x_i that satisfies,

$$\min_{x_i^{'} \in x_0 + K_i} \|b - \mathrm{A}x_i^{'}\|, \tag{7.3}$$

where $\|.\|$ is the 2-norm.

GMRES selects its first search direction v_1 to be the initial residual vector r_0. This defines the Krylov search space K_m explicitly in terms of the initial residual r_0 and matrix A,

$$K_m(\mathrm{A}, r_0) = \text{span}\{r_0, \mathrm{A}r_0, \mathrm{A}^2r_0, \ldots, \mathrm{A}^{m-1}r_0\}. \tag{7.4}$$

The algorithm generates an orthonormal basis of the Krylov subspace (7.4). At each iteration, a search direction v_i is generated through the matrix vector multiplication $\mathrm{A}v_{i-1} = v_i$ and orthonormalised with respect to the previous search directions, $v_1, v_2 \ldots v_{i-1}$. Orthonormalisation is performed using the Gram-Schmit method. The elements of an upper Hessenberg matrix H are also computed [163]. Two properties of H are,

$$\mathrm{A}V_m = V_{m+1}\hat{H}_m \tag{7.5}$$

and

$$V_m^T \mathrm{A} V_m = H_m, \tag{7.6}$$

for which the proof can be found in [154]. The matrix V_m is the $n \times m$ matrix formed through the augmentation of the column vectors of search directions $v_1, v_2 \ldots v_m$. The matrix H_m is formed from the Hessenberg matrix $\hat{H}_m$ with the last row removed. The approximation x_m is expressed in terms of the initial guess x_0 and search directions V_m by,

$$x_m = x_0 + V_m y, \tag{7.7}$$

where y is a vector of length m. The residual is expressed as,

$$r_m = b - \mathrm{A}x_m = b - \mathrm{A}(x_0 + V_m y) = r_0 - \mathrm{A}V_m y = \beta V_{m+1} e_1 - V_{m+1}\hat{H}_m y, \quad (7.8)$$

where e_1 is the first column vector of the $(m+1) \times (m+1)$ identity matrix. Using (7.8), the least squares problem reduces to:

$$\min_y \|\beta e_1 - \hat{H}_m y\|. \quad (7.9)$$

GMRES has gained much popularity for solving linear systems generated through discretising PDEs. This is due to the computationally expensive part of the algorithm being the matrix vector multiplication. Therefore, performing this operation using large sparse linear systems generated by many discretised PDEs is relatively inexpensive. However, the method does have its draw backs. Computational round off errors can dominate the solutions for large Krylov sets. The expense of storing search vectors and computing their orthonormalisation also grows linearly with each iteration of the algorithm. The method of restart is normally used to overcome these issues. This allows the algorithm to iterate between $m = 5$ and 40 iterations. The memory for the search directions is then cleared and GMRES is restarted with the latest approximation, $x_0 = x_m$.

Preconditioning

Krylov solvers can suffer with a lack of robustness resulting in a slowly converging iterative process. This problem often occurs when poorly conditioned matrices are involved. However, an effective method of avoiding this situation is to use preconditioning. The method of preconditioning re-casts the system of equations into a better conditioned linear system that possess the same solution. The preconditioner is defined by a mapping matrix M that approximates the matrix A to some extent. The method can be applied to the linear system by two approaches. One direction is the method of left preconditioning, $M^{-1}\mathrm{A}x = M^{-1}b$, and the other right preconditioning,

$$\mathrm{A}M^{-1}u = b \quad (7.10)$$

where,

$$x = M^{-1}u. \quad (7.11)$$

In both cases, the matrix systems $\mathrm{A}M^{-1}$ and $M^{-1}\mathrm{A}$ are expected to be better conditioned than the linear system A.

Right preconditioning is employed within the GMRES solver. The algorithm is modified to solve the system of equations (7.10) and then the u variables are substituted into (7.11) to obtain x. The u variables are not formed explicitly as the initial direction v_1 is calculated from the initial vector x_0, by $v_1 = \mathrm{A}M^{-1}u_0 = \mathrm{A}x_o$, and the approximation is expressed in terms of the x_m by pre-multiplying u_m by M^{-1}, to give,

$$M^{-1}u_m = M^{-1}u_0 + M^{-1}Vy = x_m = x_0 + M^{-1}Vy. \tag{7.12}$$

The preconditioned GMRES finds the approximation x_m by a projection onto the preconditioned Krylov subspace,

$$K_m(\mathrm{A}M^{-1}, r_0) = \text{span}\{M^{-1}r_0, \mathrm{A}M^{-1}r_0, (\mathrm{A}M^{-1})^2 r_0, \ldots, (\mathrm{A}M^{-1})^{m-1}r_0\}. \tag{7.13}$$

Provided M is a close approximation to A, the number of iterations required to solve the system will be reduced significantly. However, in order to be sure that the computation of $\mathrm{A}M^{-1}v_j$ does not swamp the efficiencies gained through reducing the iterations, M must be carefully selected to ensure that it is computationally efficient to solve.

FGMRES

The FGMRES method has the option of applying a different preconditioner for every iteration of the GMRES algorithm. In the FGMRES algorithm the vectors $z_j = M_j^{-1}v_j$ are defined using a unique preconditioning matrix M_j for each iteration j. The search vectors v_j are then formed by the matrix vector multiplication $\mathrm{A}M^{-1}v_j = \mathrm{A}z_j$ and orthogonalised with the previous search directions.

The approximation x_m is expressed in terms of the initial vector x_0 and search directions z_j, for $j \in \{1, 2, \ldots, m\}$. That is,

$$x_m = x_0 + Z_m y, \tag{7.14}$$

where Z_m is the $n \times m$ matrix containing the columns of search directions z_j for $j \in \{1, \ldots, m\}$. The algorithm is presented in algorithm 1. Note that line 4 can act independently of the preconditioning matrix and may therefore take the form of an arbitrary vector. This gives the algorithm complete freedom in choosing its search directions. However, the costs are that the directions Z_m are required to be stored in memory - this doubles the algorithm's storage requirements.

The least squares condition on line 14 follows from the condition,

$$\mathrm{A}Z_m = V_{m+1}\hat{H}_m. \tag{7.15}$$

This leads into expressing the residual as,

$$r \;=\; b - \mathrm{A}x \;=\; r_0 - \mathrm{A}Z_m y \;=\; \beta v_1 - V_{m+1}\hat{H}_m y \;=\; V_{m+1}(\beta e_1 - \hat{H}_m y),$$

and reducing the least squares problem to the following,

$$\min_y \|\beta e_1 - \hat{H}_m y\|. \tag{7.16}$$

Intelligent Direction Replacement: IVOR

A new method is presented in this section that provides an alternative approach to automatically restarting FGMRES once the available search direction memory has been exhausted by the algorithm. The method (employed by IVOR) continues the iterations of the solver when the search direction space is full. For each additional iteration, a search direction is discarded and replaced with a new search direction generated using the most recent residual. The new search vector is then orthonormalised with the existing $m - 1$ search directions and its contribution to the approximation is computed. The search direction discarded from the search direction set is that with largest contribution to the approximation, that is, the direction corresponding to the largest y_i in line 14 of algorithm 1 (note that the contributions are weighted with the current residual in order to make the comparison fair). The purpose for this is that by discarding the largest contributing vector, a similar search direction may be explored - since the new search direction will not

Algorithm 1 FGMRES

$r_0 = b - \mathrm{A}x_0, \beta = \|r_0\|$ and $v_1 = r_0/\beta$

Define: $\hat{H}$ as an $m{+}1 \times m$ matrix with elements $\{h_{i,j}\}$ for $1 \leq i \leq m{+}1, 1 \leq j \leq m$.

for $j = 1, m$ **do**

  $z_j = M_j^{-1} v_j$

  $w = \mathrm{A}z_j$

  for $i = 1, j$ **do**

    $h_{i,j} = (w, v_i)$

    $w = w - h_{i,j} v_i$

  end for

  $h_{j+1,j} = \|w\|$

  $v_{j+1} = w/h_{j+1,j}$

end for

Define: $Z_M = [z_1, \ldots, z_m]$

Calculate y_m through minimising $\|\beta e_1 - \hat{H} y_m\|$

$x_m = x_0 + Z_m y_m$

have to be orthogonalised w.r.t. the discarded vector. This is expected to be advantageous as the direction discarded has been shown to be a good direction to explore since it has the largest contribution of all the search directions. Furthermore, the new direction will be orthonormalised against the remaining directions that have contributed the least to the approximation. The new direction will therefore be forced to search away from these poor directions. The process of direction replacement is iterated ζ times (determined by the user) after which the algorithm is restarted with the latest approximation. The IVOR algorithm is presented in algorithm 2 and seeks an approximation x_ζ spanning the search direction space defined by:

$$K_\zeta = \{M_0^{-1}r_0, M_1^{-1}r_1, M_2^{-1}r_2, \ldots, M_\zeta^{-1}r_\zeta\}. \tag{7.17}$$

Standard Preconditioners: Stationary Iterative Solvers

Preconditioned systems are derived from the process of any stationary iterative solver of linear systems (7.1) being expressed in the form, $Bx_{i+1} = Cx_i + b$, where B and C denote an arbitrary splitting of the matrix A, $\mathrm{A} = B - C$. In [154] it is shown that this splitting leads to the preconditioned system, $M^{-1}\mathrm{A}x = M^{-1}b$, where M is the preconditioning matrix. Standard iterative solvers such as Jacobi, Gauss Seidel, SOR and SSOR are commonly employed as preconditioners. Their preconditioning matrices are expressed as, $M = D, M = D - L, M = \frac{1}{\omega}(D - \omega L)$ and

$$M = \frac{1}{\omega(2-\omega)}(D - \omega L)D^{-1}(D - \omega U), \tag{7.18}$$

respectively. The matrices D, $-L$ and $-U$ denote the diagonal, lower triangle and upper triangle of the matrix A, respectively.

7.2.2 Multigrid Scaling

This section presents the multigrid method and its application to preconditioning Krylov subspace solvers. The processes are described in the context of obtaining the solution to a linear system that is generated through the discretisation of some general PDE.

Algorithm 2 IVOR

$k = 1, r = b - \mathrm{A}x_0$.
for $j = 1, \zeta$ **do**
  $z_k = M_j^{-1} r$
  $q_k = A z_k$
  for $i = 1, \min(j, m)$ **do**
    if $(k \neq i)$ **then**
      $q_k = q_k - (q_k, q_i) q_i$
      $z_k = z_k - (q_k, q_i) z_i$
    end if
  end for
  $q_k = \frac{q_k}{||q_k||}$
  $z_k = \frac{z_k}{||q_k||}$
  $\alpha_k = (q_k, r)$
  $x = x + \alpha_k z_k$
  $r = r - \alpha_k q_k$
  $\alpha_k = \frac{\alpha_k}{\|r\|}$
  if $(j > m)$ **then**
    $k = l \quad s.t. \quad \alpha_l = \max\{\alpha_i, i \in \{1, 2, \ldots, m\}\}$
  else
    $k = k + 1$
  end if
end for

Discretising PDEs using approximations such as finite elements and finite difference generates linear systems that are to be solved in order to find the approximation's coefficients. Assuming that the element (or step) sizes of the discretisation method can be generally represented by the value h_1, the resulting linear system generated by the approximation is given by,

$$\mathrm{A}_{h_1} x_{h_1} = b_{h_1}, \tag{7.19}$$

where x_{h_1} denotes the exact solution to the linear system A_{h_1} and b_{h_1} is the corresponding right hand side vector. The error of an approximation u_{h_1} is defined as,

$$e_{h_1} = x_{h_1} - u_{h_1}, \tag{7.20}$$

and the residual r_{h_1} is given by,

$$r_{h_1} = b_{h_1} - \mathrm{A}u_{h_1} = \mathrm{A}e_{h_1}. \tag{7.21}$$

Relaxation schemes, such as Jacobi or Gauss Seidel, can be employed to relax the error terms of the approximation. However, a multigrid method functions by re-modelling the PDE over a coarser grid. The re-modelling is performed with discretisation elements (or steps) of the general magnitude $h_2 > h_1$, and produces a second, smaller linear system A_{h_2}. The residual r_{h_2} is then projected onto the coarser grid using a restriction operator $\mathcal{Q}$,

$$\mathcal{Q} r_{h_1} = r_{h_2}, \tag{7.22}$$

and the new system solved to obtain the error vector e_{h_2} defined over the coarse grid,

$$\mathrm{A}_{h_2} e_{h_2} = r_{h_2}. \tag{7.23}$$

By construction, the linear system A_{h_2} is smaller than the system A_{h_1}, and can be solved exactly to obtain the error e_{h_2}. The error e_{h_2} is then mapped to the finer grid using a prolongation operator $\mathcal{R}$,

$$\mathcal{R} e_{h_2} = \hat{e}_{h_1}, \tag{7.24}$$

resulting in an approximation to the error term e_{h_1}. The updates to the current solution u_{h_1} is then given by,

$$u_{h_1}^{new} = u_{h_1} + \hat{e}_{h_1}. \tag{7.25}$$

The procedure just described forms the basis of the multigrid method. In general, it performs well in terms of giving fast convergence to the smooth proportion of the solution. That is, the solution's general shape represented by the coarse discretisation. However, on its own the method is poor in converging on the solution's local detail represented within the finer grid. It is therefore mixed with Jacobi or Gauss Seidel relaxation, as these methods are recognised for providing quick convergence in the solution's fine detail - the combination of the two methods should therefore provide fast convergence to all aspects of the solution. The sweeps are applied across the fine grid prior to, and after performing the operations of (7.25) .

In this example two discretisations of varying resolution levels were used to create two linear systems for approximating the PDE. However, an arbitrary number of grids may be used in the general multigrid framework. This allows a fine grid to be successively coarsened an arbitrary number of times until a low resolution, easily solvable linear system is obtained. The multigrid procedure then visits and relaxes the errors over the highest and intermediate levels and solves the errors exactly for visits to the lowest resolution linear system. Using multiple levels allows a numerous number of cycles (the sequence of visiting the levels) to be employed. A number of the popular approaches are discussed in a later section.

The previous section showed how standard iterative solvers were employed as preconditioners for the Krylov subspace solvers. The preconditioning involved a block Gauss-Seidel relaxation sweep (BSOR) or a forward and backward relaxation sweep (BSSOR) through the system's elements. In the same manor, multigrid can be employed as a preconditioner by using some multigrid cycle to act as the approximation M to the linear system A in algorithms 1 and 2.

7.2.3 Streamline Upwind Petrov Galerkin discretisation of the BTE

This section first details the Streamline Upwind Petrov Galerkin FE discretisation of the angular discretised Boltzmann transport equation. This is then followed with a review of the approximation method's properties (in angle and space) for increasing the efficiency

in computing the operations of the IVOR (and FGMRES) algorithm.

This section begins with the angularly discretised steady state Boltzmann transport equation.

$$A_x \frac{\partial \Psi(r)}{\partial x} + A_y \frac{\partial \Psi(r)}{\partial y} + A_z \frac{\partial \Psi(r)}{\partial z} + \mathcal{H}\Psi(r) = \mathcal{S}(r) \tag{7.26}$$

This is generated using the Bubnov Galerkin weighted residual method, formed from the first order Boltzmann transport equation, where the angular flux $\psi(r, \hat{\Omega})$ is represented by the set of $\mathcal{M}$ angular basis functions $\mathcal{G}_i$ for $i \in \{1, \ldots, \mathcal{M}\}$.

$$\psi(r, \hat{\Omega}) \approx \sum_{j=1}^{\mathcal{M}} \Psi_j(r) \mathcal{G}_j(\hat{\Omega}) \tag{7.27}$$

The variable Ψ represents the vector of size $\mathcal{M}$ containing the angular moments corresponding to the angular basis functions $\mathcal{G}_i$ for $i \in \{1, \ldots, \mathcal{M}\}$. The term A is the $\mathcal{M} \times \mathcal{M}$ angular mass matrix. The matrices A_x, A_y and A_z are $\mathcal{M} \times \mathcal{M}$ angular Jacobians. The $\mathcal{M} \times \mathcal{M}$ matrix $\mathcal{H}$ contains the angular discretised scattering and removal operator and $\mathcal{S}$ is the vector of size $\mathcal{M}$ containing the angular discretised source.

Prior to discretising the space and time variables the angularly discretised equation (7.26) is pre-multiplied by the Petrov Galerkin term $(I - A \cdot \nabla P)$. This gives,

$$(I - A \cdot \nabla P)\mathcal{R}(\Psi(r)), \tag{7.28}$$

where

$$\mathcal{R}(\Psi(r)) = \boldsymbol{A} \cdot \nabla \Psi(r) + \mathcal{H}\Psi(r) - \mathcal{S}(r). \tag{7.29}$$

The term $\boldsymbol{A}$ represents the vector of angular Jacobians (A_x, A_y, A_z) and $\nabla = (\frac{\partial}{\partial x} I, \frac{\partial}{\partial x} I, \frac{\partial}{\partial y} I)$ where I represents the $\mathcal{M} \times \mathcal{M}$ identity matrix. The $\mathcal{M} \times \mathcal{M}$ matrix P represents the stabilising term. The role of this matrix is to add the correct amount of diffusion in the stream line direction in order to mitigate the Gibbs oscillations that form in the approximated solution. Various definitions for this matrix exists, see [19] for details. In this chapter the stabilising matrix defined in chapter 6 and [4] is used, this is defined as,

$$P'_{i,j} = \begin{cases} min((H'_{i,j})^{-1}, \frac{h}{2} A'_{i,j}) & \text{if } i = j \\ 0 & \text{otherwise,} \end{cases}$$

where,

$$H'_{i,j} = \begin{cases} \sum_k |H_{i,k}| & \text{if } i = j \\ 0 & \text{otherwise,} \end{cases}$$

and

$$A'_{i,j} = \begin{cases} \frac{1}{A_{i,j}} & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

The choice for this matrix lies with it being diagonal. It will shortly by shown that this increases computational efficiency through induced sparsity within the linear system.

The spatial variables are approximated using finite element functions. The angular flux, external source and scattering and removal operator matrix are represented by,

$$\Psi(r) \approx \boldsymbol{\Psi}(r) = \sum_{j=1}^{\mathcal{N}} \boldsymbol{N}_j(r)\boldsymbol{\Psi}_j, \tag{7.30}$$

$$\mathcal{H}(r) \approx H(r) = \sum_{j=1}^{\mathcal{N}_\mathcal{H}} \boldsymbol{N}_{Hj}(r)H_j, \tag{7.31}$$

and

$$\mathcal{S}(r) \approx S(r) = \sum_{j=1}^{\mathcal{N}_\mathcal{S}} \boldsymbol{N}_{Sj}(r)S_j, \tag{7.32}$$

respectively. The angular flux space variables are represented by the finite element functions $N_j(r)$ for $j \in \{1, \ldots, \mathcal{N}\}$. Similarly, the scattering-removal and source terms are represented by the finite element functions $N_{Hj}(r)$ and $N_{Sj}(r)$, respectively. The terms $\boldsymbol{N}_j, \boldsymbol{N}_{Hj}$ and $\boldsymbol{N}_{Sj}$ represents the $\mathcal{M} \times \mathcal{M}$ matrices containing the finite element functions N_j, N_{Hj} and N_{Sj} along their diagonal. The terms $\boldsymbol{\Psi}_j$ and S_j are vectors of size $\mathcal{M}$ containing the coefficients of the angular flux and source, respectively, associated with finite element function N_j and N_{Sj}. Finally H_j represents the scattering and removal operator matrix associated with node j of the scattering removal finite element discretisation. In this paper the condition $N_j = N_{Hj}$ = N_{Sj} is assumed for $j \in \{1, 2, \ldots, \mathcal{N}\}$.

The weak form equations are formed through pre-multiplying equation (7.28) by the $\mathcal{M} \times \mathcal{M}$ matrix $\boldsymbol{N}_j$ and integrating over space. The resulting equations can be expressed in the following form where Green's theorem has been applied to reduce some of the volume

integrals to integrals over the domain's surface.

$$\begin{aligned} -\int_V (\boldsymbol{A}\cdot\nabla \boldsymbol{N}_i(r))\Psi(r)dV &+ \int_V \boldsymbol{N}_i(r)(H(r)\Psi(r)-S(r))dV \\ +\int_V \boldsymbol{A}\cdot\nabla \boldsymbol{N}_i(r)P\mathcal{R}(\Psi(r))dV &+ \int_\Gamma \boldsymbol{N}_i(r)(\boldsymbol{A}\cdot\boldsymbol{n})(\Psi(r))d\Gamma \\ -\int_\Gamma \boldsymbol{N}_i(r)(\boldsymbol{A}\cdot\boldsymbol{n})P\mathcal{R}(\Psi(r))d\Gamma &= 0, \qquad \forall\, i \in \{1,\dots\mathcal{M}\} \end{aligned} \tag{7.33}$$

The vector n denotes the outward direction normal to the problems surface Γ. The term $\boldsymbol{A}\cdot n = n_xA_x + n_yA_y + n_zA_z$. In this formulation the last term of equation (7.33) may be presumed zero.

Efficiently Solving the BTE with IVOR

Efficiently computing Arnoldi solvers relies primarily on the discretisation methods producing sparse matrices when large linear systems are involved. Finite element methods are recognised for this property in producing sparse systems of equations when employed to discretise the spatial and time dimensions of PDEs. In relation to solving the BTE, the method of angular discretisation also plays a significant role in the solver's computational efficiency. The method chosen to discretise the angular variable determines the structure of the $\mathcal{M}\times\mathcal{M}$ block matrices generated at each pair of finite element nodes (i,j), for $i,j\in\{1,\dots\mathcal{N}\}$, that share an element. In order to compute the solvers operations efficiently, sparse block matrices are required. Each block is generated by a combination of 17 matrices defined in (7.33). Five of the 17 constituent matrices are the angular matrices of the discretised transport equation (7.26),

$$A, \qquad A_x, \qquad A_y, \qquad A_z, \qquad H. \tag{7.34}$$

Three are formed through multiplying the angular Jacobians with PH,

$$A_xPH, \qquad A_yPH, \qquad A_zPH, \tag{7.35}$$

and the remaining 9 are formed through a combination the angular Jacobians and P,

$$A_xPA_x,\ A_xPA_y,\ A_xPA_z,\ A_yPA_x,\ A_yPA_y,\ A_yPA_z,\ A_zPA_x,\ A_zPA_y,\ A_zPA_z. \tag{7.36}$$

Sparsity structure of 2-D angular matrices					
	P_9 (55)	P_{17} (171)	LW_2 (60)	LW_3 (180)	LW_4 (616)
A	1.8	0.6	25.0	20.7	11.8
A_x	5.4	1.9	25.0	20.7	11.8
A_y	5.4	1.9	25.0	20.7	11.8
H	1.8	0.6	25.0	20.7	11.8
A_xPH	5.4	1.9	25.0	24.5	24.5
A_yPH	5.4	1.9	25.0	24.5	24.5
A_xPA_x	10.3	4.1	24.5	23.9	24.0
A_xPA_y	7.7	3.0	25.0	23.5	23.8
A_yPA_x	7.7	3.0	25.0	23.5	23.8
A_yPA_y	10.3	4.1	24.5	23.9	24.0

Table 7.1: This table lists the percentage of nonzero elements of the 2-D angular matrices for the P_9, P_{17} spherical harmonic and LW_2, LW_3, LW_4 spherical wavelet angular discretisations.

In order that the block matrices retain a high degree of sparsity, the matrices defined in (7.34),(7.35) and (7.36) should also be sparse. The angular discretisation must therefore be relied upon to generate sparse angular matrices in (7.34). The stabilising matrix P was also defined in the previous section to be diagonal. This can therefore only act to preserve sparsity within the matrices (7.35) and (7.36), once the angular discretisation has been selected. Figures 7.1 and 7.2 show the sparsity structure of some of these constituent matrices generated by a P_9 and LW_2 angular discretisation and diagonal stabilising term P. Table 7.1 lists the percentage of non zero elements in these matrices using various spherical harmonic and spherical wavelet discretisations. These figures show that the discretisations go some way to producing the sparse structure required within the block matrices.

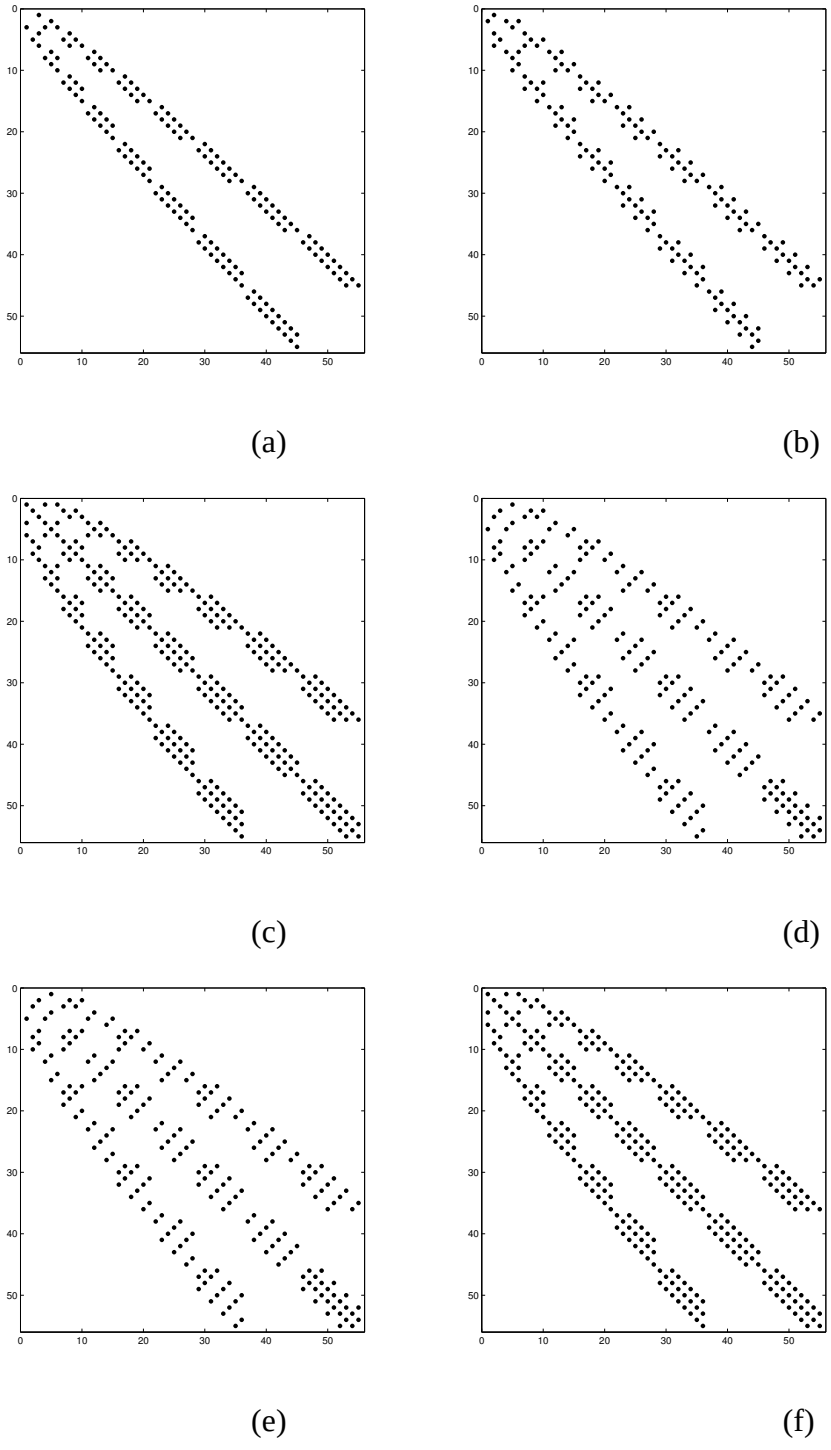

Figure 7.1: These illustrate the sparsity structure of the 2-D angular matrices generated by the P_9 angular discretisation: figure (a) A_x and A_xPH, figure (b) A_y and A_yPH, figure (c) A_xPA_x, figure (d) A_xPA_y, figure (e) A_yPA_x, figure (f) A_yPA_y

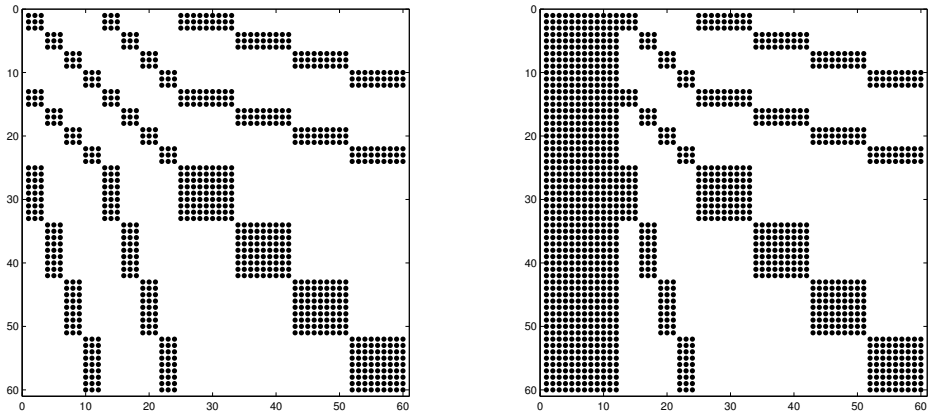

Figure 7.2: These illustrate the sparsity structure of the 2-D angular matrices generated by the LW_2 spherical wavelet angular discretisation. The left diagram shows the structure of all the matrices A_x,A_y, A_xPH, A_yPH, A_xPA_x, A_xPA_y, A_yPA_x and A_yPA_y for pure absorbing problems. The right diagram shows the structure of the matrices A_xPH and A_yPH for isotropic scattering problems.

7.2.4 Hierarchical Preconditioner

The preconditioner employed within FGMRES will play the pivotal role for determining the solver's success in converging to the exact solution efficiently. This section develops an angular multigrid preconditioner for the FGMRES algorithm that is designed specifically for solving the discretised Boltzmann transport equation. The grids (or levels) of the MG method are generated using the varying expansions of a hierarchical angular discretisation - for example spherical harmonics, Walsh functions or spherical wavelets. More specifically, the preconditioner fixes the spatial, energy and time discretisations and generates grids using the blocks of angular moments generated by each level of the angular discretisation. The multigrid cycles may then visit a level and relax or mitigate the error terms associated with the corresponding angular moments.

As with most standard multigrid methods, the error terms are relaxed when the cycle visits the highest or intermediate levels and are solved directly on a visit to the lowest order angular moments. Relaxing is performed using block SSOR (BSSOR). This involves a forward and backward sweep through the spatial discretisation's nodes. At each node a

local matrix system is formed and solved in terms of the node's angular moments associated with the level of the MG cycle. This process is analogous to the relaxed Gauss Seidel (for relaxing a single variable) extended to multi variables. Direct solving of the lowest order angular moments is performed using a separate (F)GMRES algorithm using BSSOR preconditioning.

Hierarchical angular discretisations are not only used to define the levels of the angular multigrid preconditioner but their use also simplifies the treatment of the prolongation and restriction operators used in mapping variables between the spaces of the MG levels. In order to implement the restriction process and reduce the error terms of an intermediate or the lowest MG level, the angular moments at all nodes are arranged in the order of their angular discretisation to generate the following linear system,

$$\begin{pmatrix} \mathbf{A}_{1\,1} & \mathbf{A}_{1\,2} \\ \mathbf{A}_{2\,1} & \mathbf{A}_{2\,2} \end{pmatrix} \begin{pmatrix} \Psi_1 \\ \Psi_2 \end{pmatrix} = \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{pmatrix}.$$

The vectors Ψ_1 and Ψ_2 form the complete set of angular basis coefficients. The Ψ_1 coefficients correspond to a restricted level for some multigrid cycle and Ψ_2 contain the remaining higher level moments. The sub matrices $\mathbf{A}_{1\,1}, \mathbf{A}_{1\,2}, \mathbf{A}_{2\,1}$, and $\mathbf{A}_{2\,2}$ form the global matrix and $(b_1, b_2)^T$ is the right hand side vector containing the discretised source. For some current approximation of the angular variables $\Psi^i = (\Psi_1^i, \Psi_2^i)^T$, the residual $\mathbf{r}_1^i$ of the restricted MG level can be calculated by the following,

$$\mathbf{r}_1^i = \mathbf{b}_1 - (\mathbf{A}_{1\,1}\ ,\ \mathbf{A}_{1\,2}) \begin{pmatrix} \Psi_1 \\ \Psi_2 \end{pmatrix}. \tag{7.37}$$

Note that this is only possible due to the angular discretisation being hierarchical. The error term $\mathbf{e}_1^i$ relating to the moments of the restricted MG level is calculated by solving the reduced system,

$$\mathbf{A}_{1\,1}\, \mathbf{e}_1^i = \mathbf{r}_1^i, \tag{7.38}$$

and, again, by the angular discretisation being hierarchical, the update to the current approximation is given by,

$$\begin{pmatrix} \Psi_1^{i+1} \\ \Psi_2^{i+1} \end{pmatrix} = \begin{pmatrix} \Psi_1^i \\ \Psi_2^i \end{pmatrix} + \begin{pmatrix} \mathbf{e}_1^i \\ 0 \end{pmatrix}. \tag{7.39}$$

Alternatively one could just relax the right hand side vector, $\hat{\mathbf{b}}_1 = \mathbf{b}_1 - \mathbf{A}_{1\,2}\,\Psi_2$, and solve the reduced system to obtain the new approximation for Ψ_1^{i+1}:

$$\mathbf{A}_{1\,1}\,\Psi_1^{i+1} = \hat{\mathbf{b}}_1. \tag{7.40}$$

In order to show these two approaches are equivalent, the solution from (7.40) can be expressed in terms of an update $\hat{\mathbf{e}}_1^i$ so that $\Psi_1^{i+1} = \Psi_1^i + \hat{\mathbf{e}}_1^i$. From (7.40) the following holds,

$$\begin{aligned} \mathbf{A}_{1\,1}\,\Psi_1^{i+1} &= \hat{\mathbf{b}}_1, \\ \mathbf{A}_{1\,1}\,(\Psi_1^i + \hat{\mathbf{e}}_1^i) &= \mathbf{b}_1 - \mathbf{A}_{1\,2}\,\Psi_2, \\ \mathbf{A}_{1\,1}\,\hat{\mathbf{e}}_1^i &= \mathbf{b}_1 - \mathbf{A}_{1\,1}\Psi_1^i - \mathbf{A}_{1\,2}\,\Psi_2 = \mathbf{r}_1^i. \end{aligned} \tag{7.41}$$

That is, from (7.38), the update $\hat{\mathbf{e}}_1^i$ is equivalent to the error estimate $\mathbf{e}_1^i$.

This approach has shown that the hierarchical angular discretisation naturally blends itself into the multigrid framework. The hierarchy avoids the cumbersome computation of the prolongation and restriction operators between the multigrid levels. Another important aspect is that the reduced linear systems of the lower levels retain the structure of the global moment matrix - see equation (7.40). The method therefore avoids re-computing a new linear system at every multigrid level.

7.2.5 Angular multigrid cycles

The multigrid method provides the useful option of concentrating effort in resolving specific angular moments by the setting of its preconditioning cycle. The common cycles found in the literature are initiated from the highest multigrid level. The simplest of these is the V cycle for which successive coarsening levels are visited and corresponding error terms relaxed. When the cycle reaches the coarsest grid, the errors are solved exactly. It then proceeds back up the refining levels, terminating at the maximum grid. The V cycle is illustrated in figure 7.3a where the MG levels are related to the shells of the P_N angular discretisation. Two other common MG cycles are the W and the F cycles. The W cycle is designed to pay more attention to the lower order moments. The cycle sweeps from

the maximum to the minimum grid, as seen with the V cycle. However, before returning back to the top grid, the cycle performs a number of visits over the low order moments. An example of the W cycle is shown in figure 7.3b. The W cycle is expected to have fast convergence if the exact solution's low order P_N moments are large and high P_N moments are small - that is if the angular flux is smooth and isotropic. The F cycle is designed to work in a similar manor. As before the cycle begins at the maximum level and proceeds to visit every intermediate stage to the minimum level. The cycle then begins a sequence of mini cycles. Each mini cycle begins at the lowest level, works its way up to a higher level and then returns back to the lowest level. The mini cycles are grouped so that their mini peaks are incremented one level at a time. The cycle is terminated when the maximum level is reached - see figure 7.3c. The F cycle is designed to pay the most attention to the low order moments. However, it deviates from the W cycle by still paying attention to the other spherical harmonic shells. In fact, it is designed to slowly reduce effort in resolving each shell as their order in the angular expansion increase. Its aim is to apply the correct amount of effort to each P_N shell for resolving moments generated by expansions of typical angular fluxes. That is, the coefficients of expanded fluxes have the tendency to be large for low order functions and to decrease as the function's expansion order increases. The F cycle therefore tries to mimic this pattern and apply a proportional amount of effort to each P_N shell in relation to the magnitude of their coefficients.

Alternative cycles extend from full multigrid methods. These methods employ cycles that are initiated from the lowest multigrid level and work their way up towards the maximum level. The ncycle is one example. This works on similar principals to the F cycle where increasing mini cycles build up the general sequence. Figure 7.3d and 7.3e show the ncycle = 1 and figure 7.3f shows the ncycle = 2, the latter using by a double mini cycles before incrementing to the next level. Other full multigrid sequences include the modified V and modified W cycles. The modified V employs an up and down sweep through the multigrid levels - see figure 7.3g. The modified W is performs the same operations as the modified V but has additional sweeps through the low ordered P_N moments before and after, see figure 7.3h.

CPU preconditioning/matrix vector costs

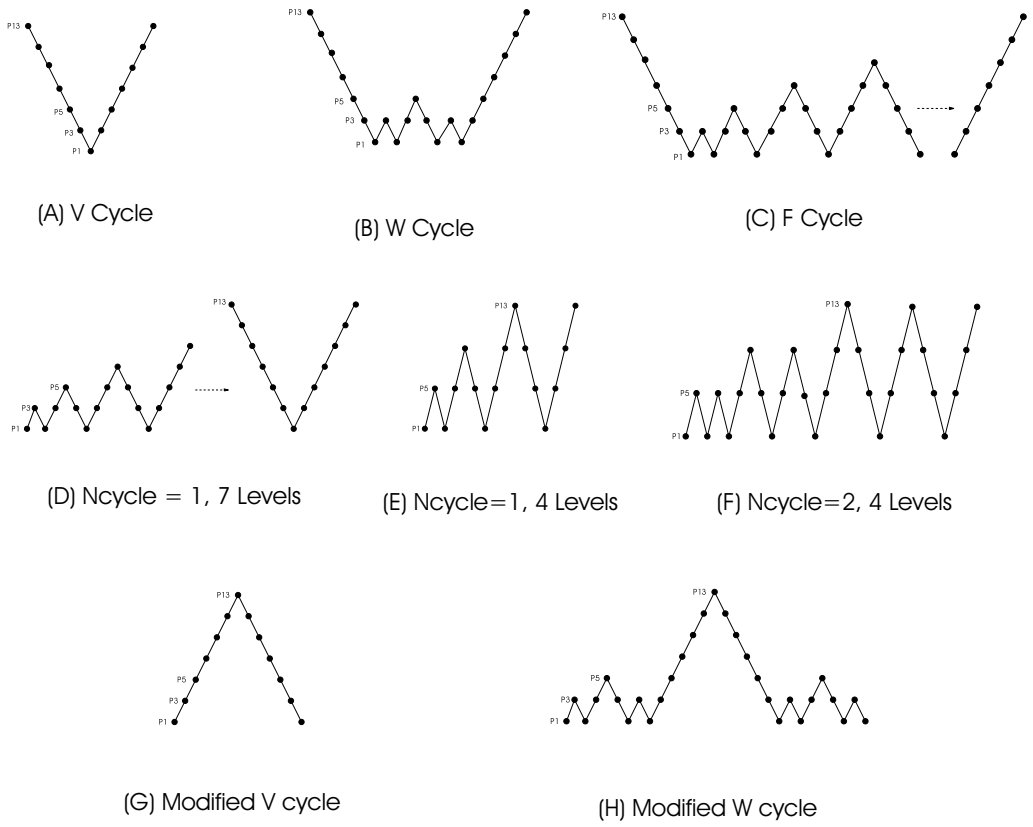

Figure 7.3: Presented here are 9 multigrid cycles over the levels generated by the P_{13} equations. Cycles A, B and C show the traditional V, W and F cycles, respectively. Cycles D, E and F show the n cycles with 4 (D and E) and 7 (F) levels. Cycles G and H show the modified V and W cycles where the cycles are initiated from the bottom level.

Preconditioning Cycle	Costs	Cost of a solver iteration
MG V	4	5
MG F	4	5
MG W	4	5
MG NCYCLE=1, 4 LEVEL	4	5
MG NCYCLE=2, 4 LEVEL	4	5
MG NCYCLE=1, 7 LEVEL	6	7
MG MODIFIED V	2	3
MG MODIFIED W	2	3
BSSOR	2	3

Table 7.2: The costs in performing the preconditioning MG cycles and an iteration of the preconditioned FGMRES and IVOR algorithms. The costs are denoted in terms of CPU units.

In order to gain an insight into the efficiency of preconditioned solver, the costs of performing the preconditioning cycles and matrix vector multiplications will need to be established. In this paper a rough estimate of the preconditioning costs are calculated by the number of visits made by the MG cycle to the highest order and most expensive angular shell (as lower MG levels are considerably cheaper to perform they may be neglected from consideration). As each visit at the top level sweeps through the angular moments twice, (a forward and backward sweep through the spatial nodes) the cost of the visit will be estimated to be two CPU units. Similarly, the matrix vector multiplication also sweeps across all spatial nodes and angular moments. The cost in performing this operation is therefore estimated to be 1 CPU unit. Table 7.2 lists the costs in performing a single iteration of the FGMRES/IVOR algorithm when preconditioned with the MG cycles illustrated in figure 7.3. The cost of the standard BSSOR preconditioner has also been included.

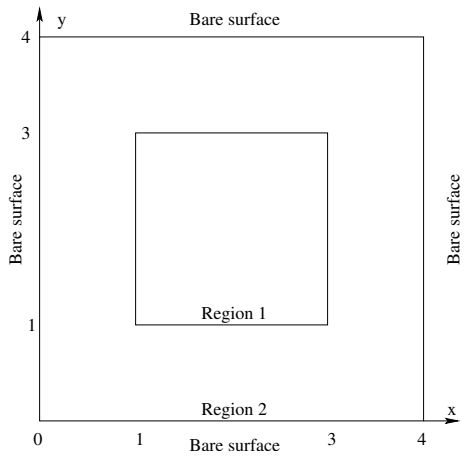

Figure 7.4: This diagram shows the domain of the scattering problem (problem 1).

7.3 Numerical Examples

This section investigates the performance of different preconditioning MG cycles for the FGMRES algorithm by solving two neutron transport problems. The domain of the first problem (problem 1) is shown in figure 7.4 and the corresponding material cross-sections are listed in table 7.3. The problem has a 4×4cm square domain containing a material with absorbing and scattering cross-sections $\sigma_a = 0.5\text{cm}^{-1}$ and $\sigma_s = 0.5\text{cm}^{-1}$. The central 2cm square region contains an isotropic source emitting 1.0 particles.$\text{cm}^{-2}\text{s}^{-1}$ and vacuum boundary conditions are prescribed across all sides. Problem 2 is a straight duct benchmark problem [49]. The geometry of the problem is shown in figure 7.6 and the corresponding material cross sections are listed in table 7.4. The problem domain has a 36×28cm rectangular region. The central 6cm square region contains an isotropic source of intensity 1.0 particles.$\text{cm}^{-2}\text{s}^{-1}$. Two ducts of width 6cm containing very low cross-sections ($\sigma_a = 0.001\text{cm}^{-1}$) extend from the source to the upper and lower boundaries. A pure absorbing material with $\sigma_a = 0.5\text{cm}^{-1}$ surrounds the ducts and vacuum boundary conditions are prescribed along all sides. The P_{13} solutions of the two problems are presented in figures 7.5 and 7.7.

The two problems used in this investigation are of varying levels of difficulty for the angular discretisation in resolving the angular flux. Problem 1 is relatively easy for the

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.5
2	0.0	0.5	0.5

Table 7.3: Source and material properties of the scattering problem (Problem 1).

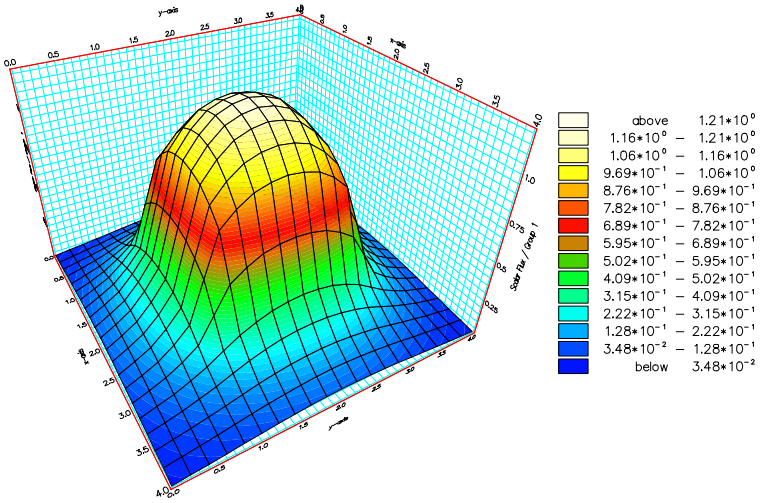

Figure 7.5: This shows the scalar flux P_{13} solution of the scattering problem (problem 1).

Region	Source $(cm^{-2}s^{-1})$	σ_a (cm^{-1})	σ_s (cm^{-1})
1	1.0	0.5	0.0
2	0.0	0.5	0.0
3	0.0	0.0	0.0

Table 7.4: Source and material properties of the Straight Duct Problem (problem 2).

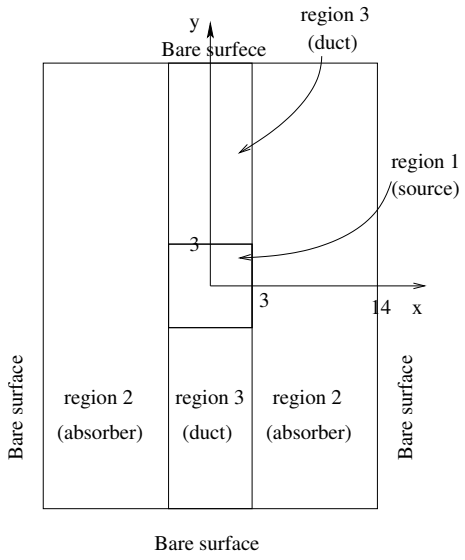

Figure 7.6: This diagram shows the domain of the straight duct problem (problem 2).

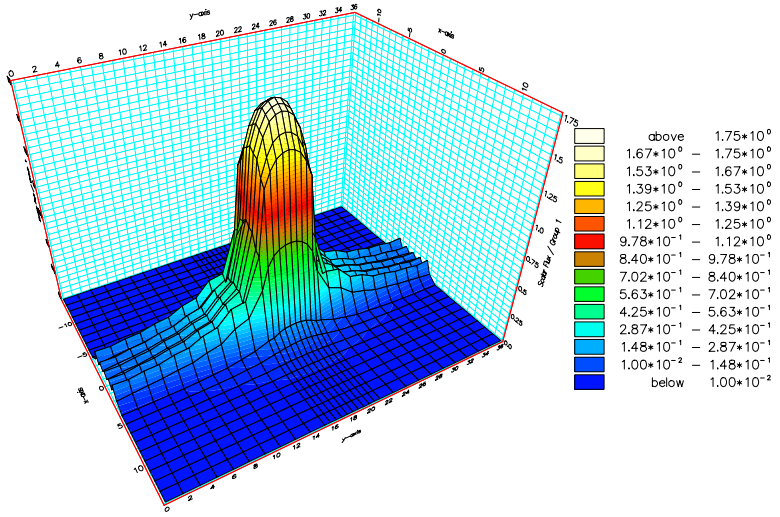

Figure 7.7: This diagram shows the scalar flux P_{13} solution of the straight duct problem (problem 2).

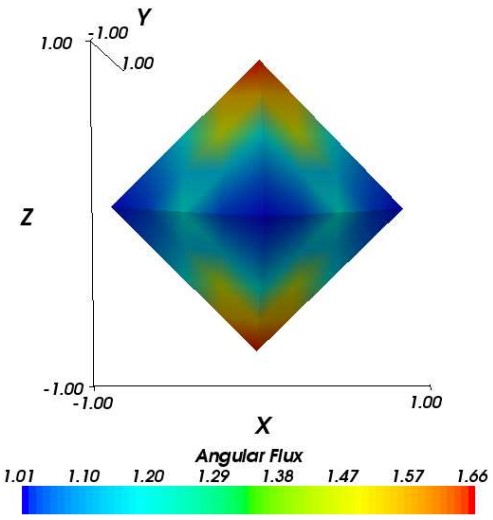

Figure 7.8: This diagram shows the angular flux LW_3 solution of problem 1 at the position (2,2). The high scattering cross-section has induced high isotropy within the angular flux. These, in turn, require only a low order angular expansion for an accurate approximation.

angular approximation method as the high scattering induces angular isotropy (see figure 7.8) which, in turn, requires only low order expansions for an accurate approximation. In contrast to this, problem 2 generates highly directed fluxes through the duct regions, an example is shown in figure 7.9. These fluxes are in general far more demanding for the angular discretisation to replicate, and in turn require large contributions from the higher order angular basis functions. The preconditioning MG cycles presented in figures 7.3 are investigated. Their performance in relation to a problem's difficulty forms the second section of this investigation. The first part aims to investigate the dependencies between the number of MG levels and the number of solver iterations. This also compares the MG performance with the standard block SSOR (BSSOR) preconditioners and also tests the effects a varying spatial discretisation has on the solver's efficiency.

The following calculations have been performed using the Streamline Upwind Petrov Galerkin code RADIANT which was partially described in the previous section. Quadrilateral finite elements are employed to discretise both problem's spatial domains - the number of elements used in the discretisation will be stated for each calculation. The hierarchical angular approximation uses both spherical harmonics and octahedral spherical

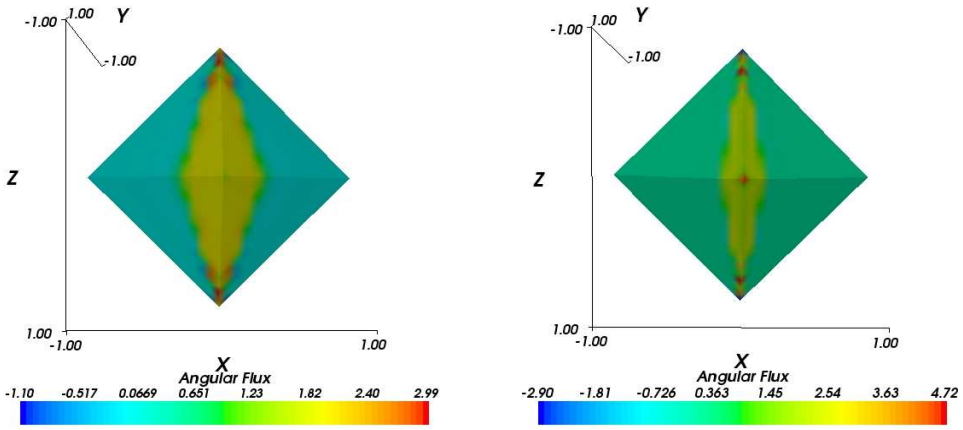

Figure 7.9: This diagram shows the angular flux LW_5 solution of problem 2 at the positions (0,27) (left) and (0,36) (right). These diagrams show the low absorbing cross-sections within the duct induce highly directed fluxes. These complex fluxes require a high order angular discretisation in order to be resolved accurately.

wavelets [1]. The multigrid preconditioning relaxing sweeps are performed using block Gauss Seidel relaxation BSSOR. The relaxation is incorporated within the $\mathcal{M} \times \mathcal{M}$ local matrix system $A^{i,i}$, (for each node i of the finite element mesh), by adding to it a diagonal matrix D^i,

$$(A^{i,i} + \frac{(1-\omega)}{\omega} D^i)\Psi_i = b^i + \frac{(1-\omega)}{\omega} D^i \Psi_{old}, \tag{7.42}$$

where D^i is formed by an absolute row sum of the elements of the matrix $A^{i,i}$,

$$D^i_{j,j} = \sum_k |A^{i,i}_{j,k}| \ , \qquad \forall j \in \{1, \ldots, \mathcal{M}\}, \tag{7.43}$$

and b^i represents the corresponding right hand side vector of the local system. The variable $\omega = 0.5$ is used to ensure all local matrices $A^{i,i}, \forall i \in \{1, \ldots \mathcal{M}\}$, is non singular and well conditioned (by it being diagonally dominant). Unless started otherwise, the following calculations have all used 20 Krylov search directions. On the event that all direction were used the method of restarting FGMRES was employed.

7.3.1 Convergence properties of MG and BSSOR preconditioners

This section is designed to test the effect the number of multigrid levels employed within the MG preconditioned FGMRES solver has on the solver's rate of convergence. For this investigation, problems 1 and 2 were solved using two spatial meshes. Meshes consisting of 200 and 400 finite elements were used to discretise the spatial domain of problem 1. For problem 2, meshes of 900 and 3,600 elements performed the spatial discretisation. For this investigation the modified V multigrid cycle was employed to solve the spherical harmonic equations, from P_3 to P_{19}, and the spherical wavelet equations, LW_1 to LW_3. Comparisons of the performance and efficiency of the MG method are made with the standard BSSOR preconditioner. For all calculations the FGMRES algorithm was allowed to run until the convergence criteria,

$$\frac{\|r\|}{\|r_0\|} \leq 10^{-5}, \tag{7.44}$$

for problem 1 and,

$$\frac{\|r\|}{\|r_0\|} \leq 5 \times 10^{-5}, \tag{7.45}$$

for problem 2, was reached. A maximum of 400 and 2000 solver iterations were allowed for the P_N and wavelet calculations.

Table 7.5 presents the number of iterations required by the FGMRES algorithm in finding the solutions of the problem 1 (P_N) when preconditioned with BSSOR and MG. The figures show FGMRES uses fewer iterations when preconditioned with MG as opposed to BSSOR. The difference in the number of iterations is at its greatest for the high order P_N calculations. Here it is seen that FGMRES used between a third and a half of the iterations when preconditioned with MG. The results also reveal the multigrid method has retained its important behaviour in that the number of solver iterations converges with respect to the number of multigrid levels. In fact, the isotropy of the problem has led to a reduction in the number of iterations for the high order P_N calculations. This is due to the high order moments being close to zero which, in turn, allows the extra relaxation sweeps of the high order P_N cycles to give further resolution to the low order P_N moments.

Similar patterns have also emerged in the number of FGMRES iterations required to solve the P_N equations generated by problem 2. These figures are listed in table 7.6. For the

Number of solver iterations using P_N on problem 1					
		Low resolution mesh		High resolution mesh	
P_N (MG levels)	Expansion size	MG	BSSOR	MG	BSSOR
P_3 (2)	10	20	23	23	32
P_5 (3)	21	20	26	22	34
P_7 (4)	36	18	27	22	37
P_9 (5)	55	15	27	21	35
P_{11} (6)	78	22	28	21	37
P_{13} (7)	105	12	28	18	38
P_{15} (8)	136	11	29	18	38
P_{17} (9)	171	11	29	18	38
P_{19} (10)	210	10	29	18	38

Table 7.5: This table displays the number of iterations performed by the FGMRES solver to resolve problem 1 discretised using various P_N angular expansions. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 10^{-5}$. The multigrid modified V cycle and BSSOR preconditioners are compared.

low resolution spatial mesh calculations, the number of FGMRES iterations converged to around 81 as the P_N discretisation was increased. This convergence pattern began to emerge in the BSSOR results, however, the number of iterations required to solve the problem was significantly higher than MG at around 250 -270. For the high spatial resolution mesh, the MG preconditioner has performed far better than BSSOR. The calculation using the BSSOR preconditioner have all exceeded the 400 iteration limit with the exception of P_3. On exceeding this limit, the relative residuals, denoted in brackets, show the BSSOR calculations to some way off from converging on its solution. The MG conditioned solver has however shown strong signs of convergence, requiring around 300 iterations when the angular expansion reached high order P_N. The wavelet calculations listed in table 7.7 also agree with the results seen so far. That is, they show the number of iterations required by the solver was significantly reduced using the MG preconditioner.

The results indicate that the efficiency of the angular multigrid preconditioned FGMRES is dependent on the resolution of the spatial grid. That is, for both problems, the number of iterations increase significantly for the refined spatial mesh. In order to investigate this further the P_5 equations were solved on a number of spatial meshes of varying resolution. For problem 1 spatial meshes consisting of 100, 400, 1,600, 6,400, 9,216 and 14,400 finite elements were used to discretise the P_5 equations and solved with the modified V preconditioned FGMRES. Similar calculations were also performed to solve problem 2 for which meshes consisting of 225, 900, 2,025, 3,600, 5,625, 10,404 and 14,400 elements were used to discretise the spatial domain. Convergence in the solutions were assumed when criteria (7.44) and (7.45) were reached for problems 1 and 2, respectively. Figure 7.10 shows the graphs that plot the number of iterations used by FGMRES to resolve the linear systems generated by all meshes. The results confirm that for both problems, the number of solver iterations increase with the number of elements used in the spatial discretisation. In fact the results show that the number of iterations increase linearly with respect to the number of spatial elements - this is especially the case for the straight duct problem.

Number of solver iterations using P_N on problem 2					
		Low resolution mesh		High resolution mesh	
P_N (MG levels)	Expansion size	MG	BSSOR	MG	BSSOR
P_3 (2)	10	65	104	219	262
P_5 (3)	21	88	135	319	-(7.0E-5)
P_7 (4)	36	92	164	362	-(1.4E-4)
P_9 (5)	55	91	195	338	-(1.3E-4)
P_{11} (6)	78	89	220	302	-(1.6E-4)
P_{13} (7)	105	88	328	287	-(1.7E-4)
P_{15} (8)	136	84	253	266	-(2.1E-4)
P_{17} (9)	171	81	254	317	-(1.7E-4)
P_{19}(10)	210	81	263	339	-(1.6E-4)

Table 7.6: This table displays the number of iterations performed by the FGMRES solver to resolve problem 2. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 5 \times 10^{-5}$ or a maximum of 400 iterations is reached. In the case of the maximum iterations being reached, the (.) indicates the relative residual at this termination point. The multigrid modified V and BSSOR preconditioners are compared.

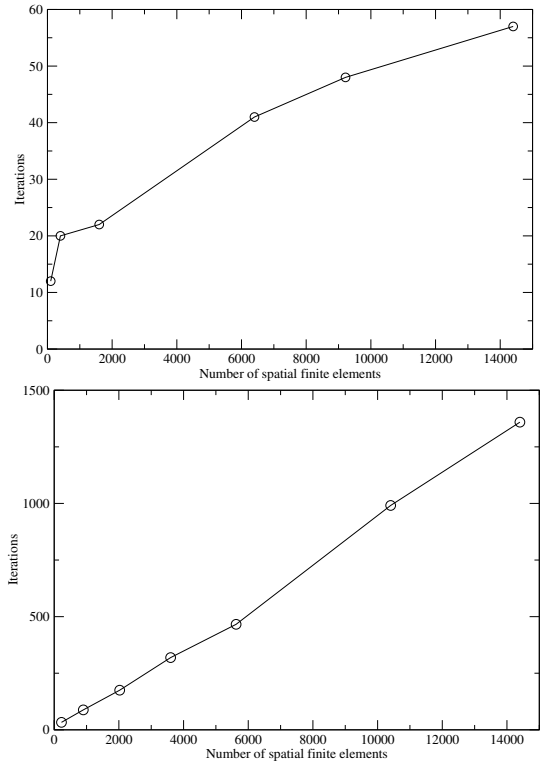

Figure 7.10: These graphs show the iterations required by the solver in resolving problem 1 (left) and problem 2 (right) on increasing resolution spatial grids. A P_5 angular discretisation is used for both problems and the solver employs a modified V MG preconditioner. The iterations are seen to increase linearly with respect to the number of spatial elements.

Number of solver iterations using LW_N on problem 2			
		Low resolution mesh	
LW_N (MG levels)	Expansion size	MG	BSSOR
LW_1 (2)	24	95	288
LW_2 (3)	60	487	1958
LW_3 (4)	180	699	-(2.3E-2)

Table 7.7: This table displays the number of iterations performed by the FGMRES solver to resolve problem 2. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 5 \times 10^{-5}$ or a maximum of 2000 iterations is reached. In the case of the maximum iterations being reached, the (.) indicates the residual at this termination point. The multigrid modified V and BSSOR preconditioners are compared.

7.3.2 Comparison of MG preconditioners

This section is designed to compare the performance of the various MG preconditioning cycles presented in figure 7.3. For this investigation problem 1 and problem 2 were solved using FE spatial meshes consisting of 200 and 900 elements, respectively, and a P_{13} angular expansion. The convergence criteria (7.44) was employed to terminate the algorithm, for which a maximum allowance of 400 iterations was made. Table 7.8 lists the number of iterations required by the solver when preconditioned using the MG cycles and BSSOR. The estimation of the CPU costs have also been included.

The results show the best cycles for reducing the number of FGMRES iterations and CPU costs for the scattering problem (problem 1) are the F, W and n cycles. These MG preconditioners led to the FGMRES processing between 5 and 8 iterations - requiring between 30 and 35 CPU units. The modified V and W cycles required 12 and 13 iterations (36 and 39 CPU units), respectively. This shows that despite giving a higher priority to the low order P_N moments, which is where the largest angular coefficients were expected to form, the modified V and W preconditioners did not perform as well as expected. The graph presented in figure 7.11 shows the residual's evolutions within the FGMRES solver. The BSSOR preconditioner is seen to perform poorly in comparison to all MG methods.

A comparison of various preconditioners		
Preconditioner	Scattering problem	Straight duct problem
MG V	7 (35)	139 (695)
MG F	6 (30)	69 (345)
MG W	7 (35)	81 (405)
MG NCYCLE=1, 4 LEVEL	8 (40)	113 (565)
MG NCYCLE=2, 4 LEVEL	6 (30)	73 (365)
MG NCYCLE=1, 7 LEVEL	5 (35)	66 (462)
MG MODIFIED V	12 (36)	140 (420)
MG MODIFIED W	13 (39)	139 (417)
BSSOR	28 (84)	397 (1191)
LUMPED DIAGONAL	65	$\sim$2000
NO	166	$\sim$1700

Table 7.8: This table displays the number of iterations performed by the FGMRES solver to resolve problems 1 and 2. The table compares the performance of all multigrid preconditioning cycles listed in figure 7.3. The table includes the figures using the standard BSSOR preconditioning, lumped diagonal preconditioning and no preconditioning. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 10^{-5}$ for both problems 1 and 2. A P_{13} angular discretisation was used to represent the angular dependence. Meshes consisting of 200 and 900 elements were used to discretise the spatial dimensions of problems 1 and 2, respectively. The (.) represents the estimated CPU costs.

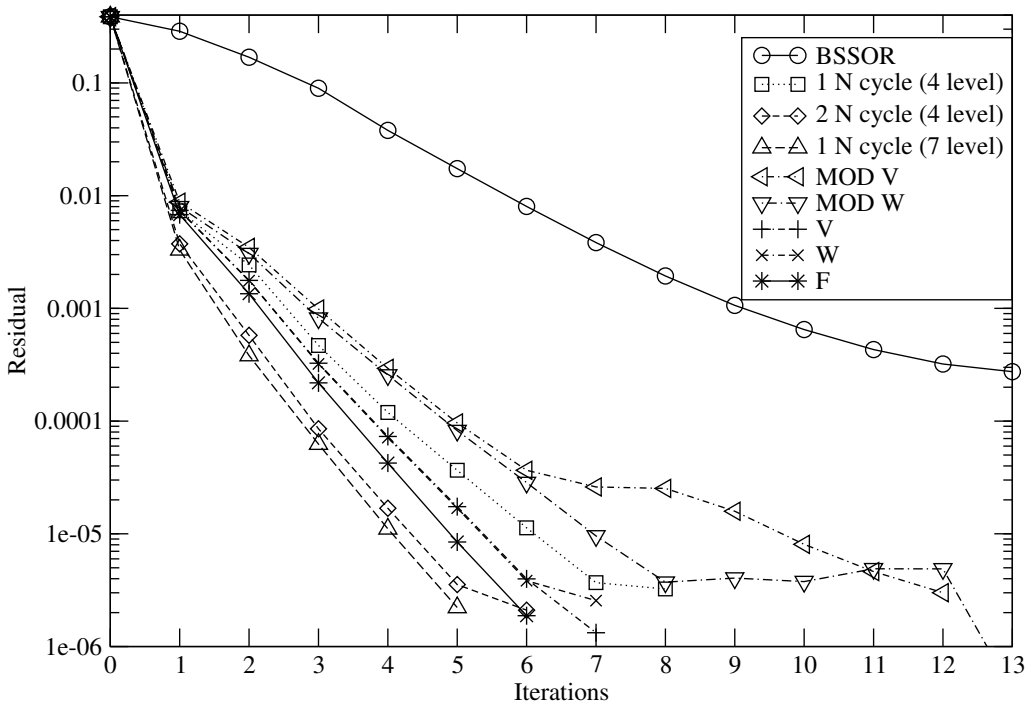

Figure 7.11: This graph presents the residual evolution when using the MG preconditioning cycles of figure 7.3 for solving the scattering problem (problem 1). Included in the graph is the residual from using the BSSOR preconditioner. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 10^{-5}$.

For all MG preconditioners, the residuals decays exponentially with respect to the number of iterations up to 7th iteration. After this point the residuals corresponding to the modified V and modified W cycles begin to level out. However, the modified W preconditioner's graph shows small increases in the residual, indicating that computational errors can dominate the solvers updates at around $r_m = 5.0 \times 10^{-6}$. Therefore, on inspection of the residual prior to the error, the modified W cycle compares well with the other preconditioners - requiring 8 iterations and only 24 CPU units - and shows the cycle to have the greatest potential with respect to this problem.

The results for the straight duct problem (problem 2) give a different insight into the quality of the MG preconditioning cycles. The figures in table 7.8 show the most effective cycles for reducing the number of solver iterations to be the n1 (7 level), followed by the F,

n2 (4 levels) and W cycles. These required 66, 69, 73 and 81 iterations respectively. When considering computation costs the F followed by the n1 (7 level) and W cycles were the most efficient. Again, the modified W and modified V cycles performed poorly, requiring 140 and 139 iterations (and over 415 CPU units), respectively. This, however, would be expected for this type of problem as these preconditioners give less preference to the high order P_N moments - thus the large coefficients corresponding to the high order moments within the ducts would be slow to evolve. The V cycle also performed poorly, requiring 140 iterations, again, for the same reasons that few visits to the high order moments were carried out. The graph in figure 7.12 shows the residual's evolution during the solver's process using all MG preconditioners. The BSSOR is again shown to have the slowest decline. The performance of the F, W and long n-cycles can be seen to be working well and produce consistently sharp reductions in the residual over all iterations. The small n cycle, modified V and modified W cycles show a residual with a more gentle decline. Note here that for these calculations, computational errors have not effected the evolution of the solver.

These results have established an insight into which preconditioners are optimal for a given problem to be solved. Although the (easy) scattering problem results revealed a smaller number of iterations were used by the F, W and ncycles, the cheaper modified V and modified W cycles required only a few more iterations. Therefore, in terms of efficiency, the modified V and, more specifically, the modified W cycles performed the best. However, in order to resolve the straight duct problem the results showed that using the more expensive F, W and n1 (7 level) cycles gave significant reductions in the number of solver iterations. The results also showed that the gains made in performing significantly less iterations compensated for the higher computational costs of the preconditioner. Therefore, when solving these type of problems it is beneficial to use the expensive preconditioners that give more preference to the high order moments.

7.3.3 IVOR's Intelligent Direction Replacement

This section is designed to demonstrate the capabilities of the direction replacement method of IVOR developed in section 7.2.1 for ensuring fast convergence when solving

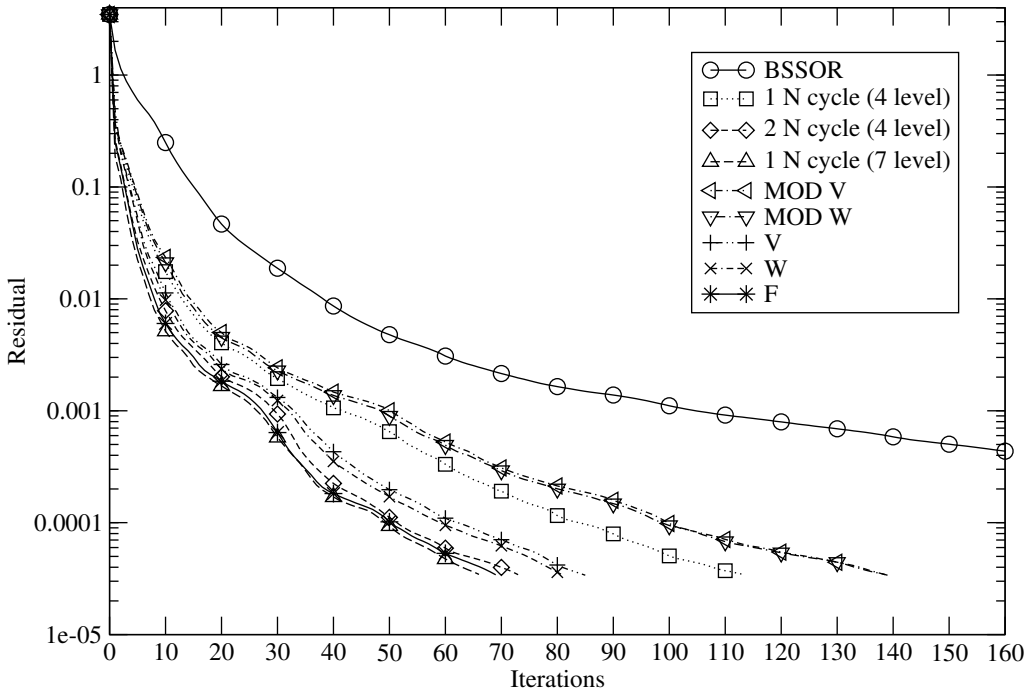

Figure 7.12: This graph presents the residual evolution when using the MG preconditioning cycles of figures 7.3 for solving the straight duct problem (problem 2). Included in the graph is the residual from using the BSSOR preconditioner. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 10^{-5}$.

difficult neutron transport problems. The investigation of this section also concentrates on the solver's convergence rates when a reduced memory for the search directions are used. This investigation solves the straight duct problem where a P_7 angular discretisation and a 14,400 finite element spatial mesh are used to discretise the angle and space dimensions. In total, six calculations have been carried out using FGMRES that is preconditioned with the standard V MG cycle. The calculations are divided equally between using 10 and 20 Krylov vectors. For both Krylov vector spaces, 3 calculations were carried out using the following options:

1. No vector updates (Standard FGMRES),
2. Vector updates allowing FGMRES to perform 100 iterations but not using latest residual to calculate the search directions,
3. Vector updates allowing FGMRES to perform 100 iterations and using latest residual to calculate the search directions (IVOR).

The graph presented in figure 7.13 shows the residuals of the six calculations. The results using 10 Krylov vectors vary greatly when the various FGMRES options 1, 2 and 3 are employed. The option using standard FGMRES produces a residual that declines slowly and requires 261 iterations to converge on the solution. The result for option 2 shows that using search vector updates decreases the number of iterations required for convergence to 207. However, the residual's profile for this calculation also reveals that during the evolution of the algorithm, the residual had the tendency to stagnate. The first of these stagnant periods occurs over the interval of 30-100 iterations. The residual only manages to free it self from stagnation when the algorithm underwent a complete restart, however, the residual stagnated again 50 iterations after this restart. The best performing variant is shown to be option 3 (IVOR). This calculation required only 124 iterations in order to find convergence. The calculation's residual graph also shows that by using the latest residual with the direction replacement method, the residual avoids all stagnation periods that were seen in the calculation of FGMRES variant 2.

The calculation corresponding to variant 1 FGMRES using 20 Krylov vectors has produced consistently sharp reductions in the residual for all iterations. The graph corre-

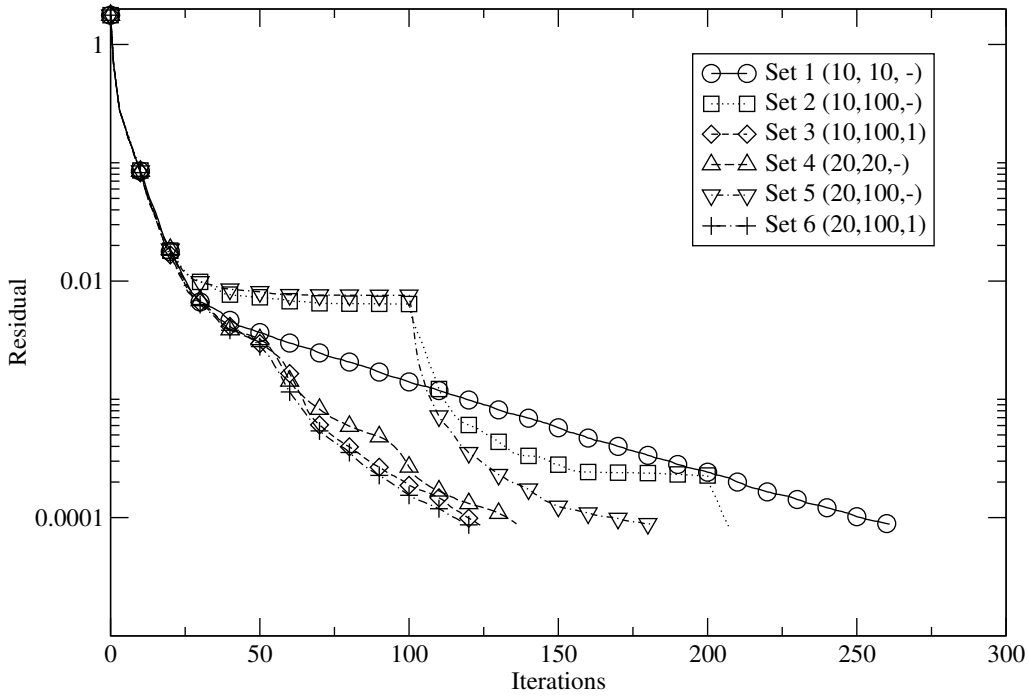

Figure 7.13: This graph presents the number of iteration required by the solver to resolve the straight duct problem using a P_7 angular expansion and a spatial mesh consisting of 14,400 elements. The results compare the reduction in the residual using the variants of the FGMRES solver discussed in section 7.2.1. Set 1: 10 Krylov vectors and restarting after 10 iterations. Set 2: 10 Krylov vectors and restarting after 100 iterations. Set 3: 10 Krylov vectors and restarting after 100 iterations using partial restart $\eta = 1$. Set 4: 20 Krylov vectors and restarting after 20 iterations. Set 5: 20 Krylov vectors and restarting after 100 iterations. Set 6: 20 Krylov vectors and restarting after 100 iterations using partial restart $\eta = 1$. All calculation were performed using the V preconditioning cycle. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 5 \times 10^{-5}$.

sponding to this calculation shows that convergence on the solution is found using just 136 iterations. Variant option 2 ,using 20 Krylov vectors, has again shown a susceptibility to a stagnating residual. As before, a complete restart was required in order to free the residual from its stagnating regions. However, in this case, the number of iterations required for convergence increased to 181. The strongest option using 20 Krylov vectors was again shown to be FGMRES variant 3 (IVOR). Option 3 consistently produced smaller residuals than both variants 1 and 2 and converged on the solution using just 121 iterations. This third variant setting has again successful avoided the stagnating residuals that occurs in the calculations using variant 2.

The results contained in figure 7.13 have shown the larger Kyrlov spaces produce faster converging solvers. This was particularly apparent in the algorithms where the standard FGMRES as used. The FGMRES variant that used the most recent residual with the direction replacement method (IVOR) performed remarkably well for both search direction sets. The calculations using this variant of FGMRES showed the size of the search spaces tested had little influence on the evolution of the residual. In fact, using the smaller search vector set required only 3 more iterations. This indicates that IVOR, using 10 search vectors may be the most efficient of all the FGMRES algorithm's variations examined here. This is based on when the reduced memory requirements and reduced computational burden of orthonormalising smaller sets of search vectors are taken into consideration.

7.3.4 Pivoting on P_1 and P_3

This section demonstrates the convergence rates of the solver when the MG preconditioner pivots on the P_3 equations, that is, where the preconditioner solves the P_3 equations directly rather than P_1. For this demonstration both the scattering and straight duct problems were solved using the P_9 and P_{11} angular approximations. The standard V MG cycles were used to precondition the solver. The preconditioner that pivots on the P_1 equations uses the standard V cycle illustrated in figure 7.3. The cycle that pivots on the P_3 equations follows the same pattern but has a turning point, where the equations are solved directly, that occurs on the P_3 shell. Figure 7.14 show the evolution of the residuals for the calculations solving problems 1 and 2.

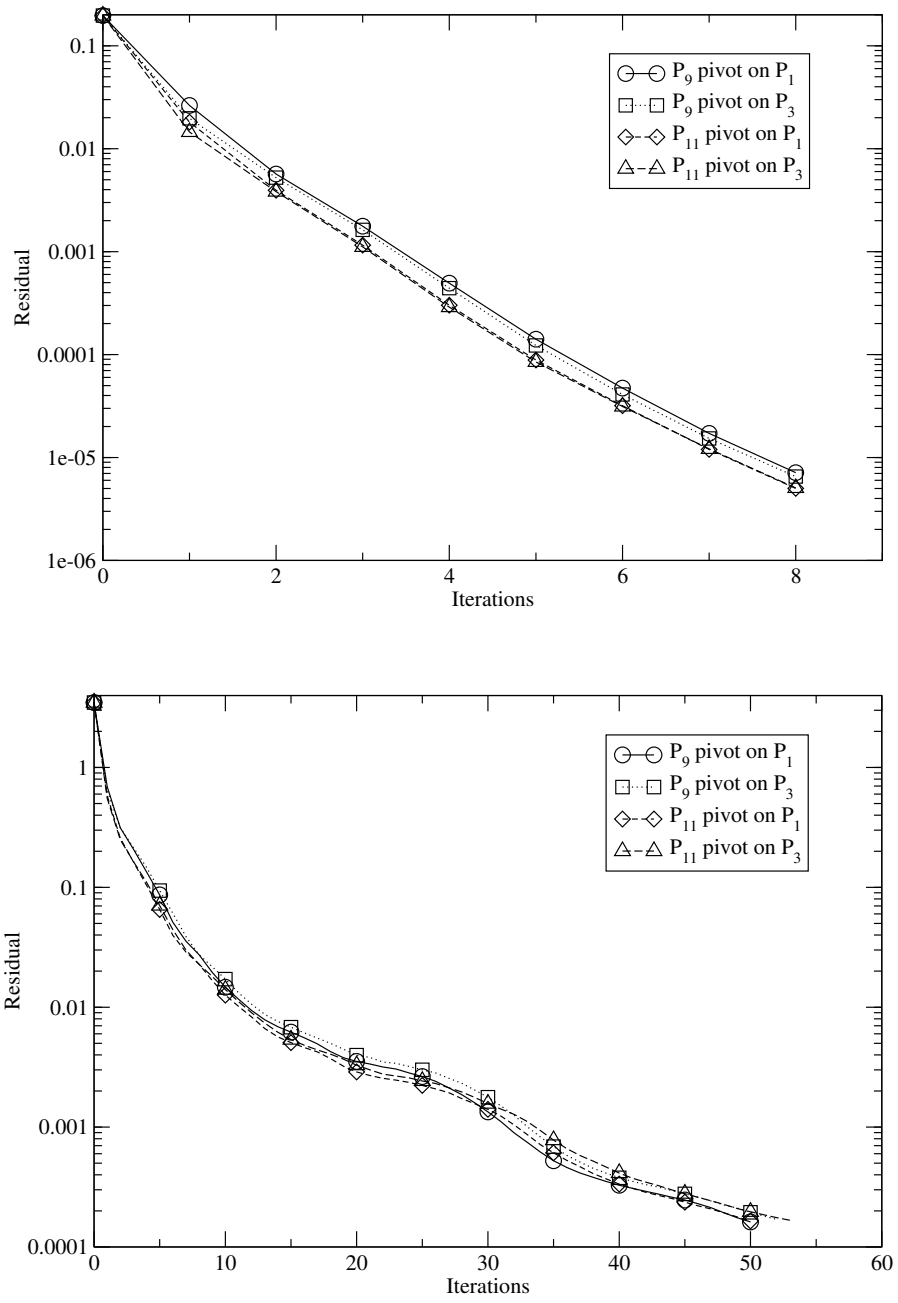

Figure 7.14: These graphs show the residual's evolution through the FGMRES solver, left: scattering problem, right: straight duct problem. The graphs compare the performance of the solver when the MG preconditioner is pivoted on the P_1 and P_3 spherical harmonic equations. The P_9 and P_{11} angular expansions are used to discretise the angular dimensions. The spatial dimensions are discretised using finite elements meshes consisting of 1,600 (scatter problem) and 900 (straight duct problem) elements. The standard V MG cycle is used to precondition the FGMRES solver. Convergence is assumed when $\frac{\|r\|}{\|r_0\|} \leq 10^{-5}$ and $\frac{\|r\|}{\|r_0\|} \leq 10^{-5}$ for problems 1 and 2, respectively.

For the scattering problem, pivoting on the P_3 equations give slight improvements in terms of reducing the residual in comparison to pivoting on P_1. The difference in the residuals between the two preconditioners are more apparent in the early iterations of the solver. However, this variation is not enough to decrease the number of iterations required by the solver when pivoting on P_3. The graph shows that all calculations required 8 iterations.

For the straight duct problem, P_3 pivoting takes very slightly longer to converge - here it is shown to be only one or two iterations more than when P_1 pivoting was used. However, this may not be too unexpected as the straight duct problem is more concerned with resolving the high order P_N moments. Therefore, solving the P_3 equations exactly will have little influence on the quality of the preconditioning.

7.4 Conclusions

In this chapter a hierarchical preconditioner for the FGMRES algorithm has been developed for solving the discretised Boltzmann transport equation. The preconditioner has been designed on algebraic and geometric multigrid methods in which hierarchical angular discretisations were used to form the multigrid levels. This chapter has also developed a new variant of FGMRES (IVOR) which offers a new alternative method to restarting when the memory for search directions is full. The preconditioned solver was incorporated within the Streamline Upwind Petrov Galerkin, finite element, radiation transport code RADIANT. A discussion on the code's efficient implementation within the algorithm has been presented in the text. Spherical wavelet LW_N and spherical harmonic P_N hierarchical angular discretisations were employed in the numerical examples. Using these discretisation methods a number of multigrid cycles were investigated.

The numerical examples have illustrated the angular multigrid method to be a good performing preconditioner for the FGMRES algorithm. They have shown that convergence in MG angular levels can be gained using even the simplest of multigrid cycles - namely the modified V cycle. The numerical examples have demonstrated that the optimal preconditioner is dependent on the problem being numerically approximated. Simple problems

with smoothly varying fluxes can be solved efficiently using cycles that concentrate on resolving low order angular moments - the modified V and modified W cycles. However, for difficult transport problems the complex preconditioning cycles that gave more preference to the high order moments were found to be the most efficient. In all examples the MG preconditioner out performed the BSSOR preconditioner by some considerable margin.

The variant of FGMRES (IVOR) developed in this chapter has been demonstrated to increase further the efficiency of the solver - this is in respect to both the CPU efficiency and the memory requirements. In particular, IVOR's method of using the latest residual to select its next search direction, together with the method of search direction replacement, was found to be significant in producing the best results on all examples tested. It was also demonstrated that the number of iterations performed by the MG preconditioned solver was dependent on the spatial resolution. That is, increases in the spatial meshes often resulted in more iterations being required by the solver. However, convergence in the solver's iterations with respect to the spatial discretisation is not expected as the MG is based on angle only. Therefore, in order to solve problems with large spatial discretisations, future research may be necessary in the mixing of this angular MG preconditioner with one for the spatial variables in order to gain iteration convergence in both space and angle.

Chapter 8

Conclusions

Contents

This section presents a conclusion of the research for neutron transport approximation that has been covered in the chapters of this thesis. Following this is a discussion of the possible directions for future work that can continue on the research presented here.

8.1 Spherical Wavelets (Chapters 3, 4 and 5)

In chapters 3,4 and 5 the foundations of the work in this thesis were laid. These chapters were concerned with the development of the three spherical wavelet bases. Low order octahedral and hexahedral spherical wavelets and a Chebyshev spectral spherical wavelet bases were developed. The wavelet bases were all applied to discretise the unit sphere, used to represent the direction of particle travel, and represent the angular variable of the angular flux of the first order Boltzmann transport equation. All bases were incorporated within the neutral particle transport code RADIANT. The code employed a Streamline Upwind Petrov Galerkin finite element method to discretise the spatial dimensions of the transport equation. A Riemann method was also used to resolve bare surface boundary conditions. This utilised the eigen-structure of the angular Jacobians in order to calculate the incoming and outgoing information through the surface of the boundary.

A number of numerical examples were presented in order to demonstrate the wavelets capabilities in approximating the Boltzmann transport equation. The wavelets abilities were clearly demonstrated to provide accurate solutions for all test problems solved. They were shown to be comparably accurate to the established methods S_N and P_N when similar sized angular expansion were used. Furthermore, in some examples the wavelet discretisations were superior to P_N in providing accurate scalar solutions. This was highlighted in the test problem involving the source within a void cylinder. This numerical example demonstrated a severe weakness of the P_N method as it was unable to generate a solution in close agreement with the exact solution. The wavelets, however, avoided this problem and did produce meaningful solutions that were in close agreement with the exact solution. The investigation of the numerical examples also set out to illustrate the wavelets ability to suppress ray-effects forming in their scalar flux solutions. This indeed turned

out to be the case as all numerical examples showed the wavelet's solutions to be less oscillatory (in their scalar flux solution due to ray effects) than the solutions of S_N. In general the numerical examples demonstrated the wavelets could produce accurate solutions for a range of extreme transport conditions, from transparent to opaque material regimes, whilst suppressing the formation of ray-effects.

Of all the wavelets developed in these chapters the linear octahedral wavelets stand out to the preferred choice for expanding in angle. The reasons for this lies with their angular matrices being the least dense. The numerical examples showed all expansions to be roughly of equal accuracy - with respect to the number of angular unknowns. Therefore, the sparser structure of the linear octahedral wavelets will serve to make the solving of their discretised equations the most efficient to perform. However, having said this, it still remains to further test these expansion on 3-dimensional problems. It is quite possible that the hexahedral based expansions are more suitable for representing these fluxes where the co-latitude angle becomes more important to representing the transport of particles.

8.2 Angular Adaptivity (Chapter 6)

In chapter 6 a wavelet method for anisotropic resolution on the sphere for discretising the angular variable of the BTE was developed. This method employed the low order octahedral and hexahedral linear wavelets, developed in chapters 3 and 4, to represent the angular flux. The method essentially detailed a technique allowing the wavelet to locate areas on the sphere in which the angular flux was under resolved and to select the correct wavelets in order to increase the angular resolution to these regions only. The chapter also discussed a routine in order for the self adaptive wavelet solution to evolve during the process of an iterative solver of the linear system generated by the space and angle discretised equations.

Several numerical examples were used to demonstrate the capabilities of the adaptive wavelets. In particular, problem domains involving ducts containing voids embedded within strong absorbers were used to generate highly demanding and highly directional

flux distributions. The angular flux solutions demonstrated the wavelets to be capable of locating the regions of angle in which the flux was directed, and to correctly apply their effort in resolving these angles only. The numerical solutions also demonstrated the wavelets to be capable of detecting regions of negligible scalar flux, and to avoid the placement of further resolution to these areas. The chapters finding showed that for practical problems such as the dog-leg duct problem, the adaptive angular resolution could reduce the number of angular moments used in the approximation by up to 80% without incurring significant errors. This in turn was shown to substantially reduce the computation times for solving the transport problems.

8.3 Fast FGMRES solver (Chapter 7)

In chapter 7 a new solver (IVOR) and preconditioner was developed in order to solve the space and angle discretised transport equations. The solver forms a new variation of the FGMRES algorithm that provides an alternative method to the restart procedure once the available memory for search directions has been exhausted. IVOR's method is to intelligently select and reject directions from the full complement of search vectors and replace them with the most recent preconditioned residual.

The preconditioner is based on a new hierarchical angular multi-grid method. The MG grids are generated using the levels of a hierarchical angular discretisation - the hierarchy is shown to naturally blend itself into a multi-grid framework. The numerous cycles available provides the solver much flexibility in selecting the effort it applies to resolve particular angular moments.

The numerical examples demonstrated the angular multi-grid preconditioner to be highly successful in increasing the efficiency of the solver in comparison with standard preconditioners. The angular multi grid cycles were also demonstrated to have a large influence on the solvers efficiency. That is, the best performing preconditioning cycles were those which suited the problem being solved. For example, isotropic problems were efficiently solved using cycles that concentrated on the low order moments whereas the problems

with highly directed fluxes were solved efficiently using cycles that concentrated on the high order angular moments. The preconditioned solver was also demonstrated to have the important property of converging in its number of iterations with respect to the size of the angular expansion. The numerical examples had also demonstrated IVOR's method of search direction replacement to increase the solver's efficiency compared to standard FGMRES.

8.4 Future Work

The research contained in this thesis has opened a number of possible avenues for future work. Some of these directions are discussed here.

8.4.1 Further Wavelet Development

The wavelets currently developed are linear, quadratic and spectral wavelets based over an octahedral or hexahedral domain [1, 2] in which all fall into the category of second generation wavelets. It therefore remains an open field for exploring the classical wavelet functions usage in neutron transport theory [144]. These types of wavelet functions offer additional advantages such as full orthogonality of high order wavelet functions. This is important as it means that the matrices have highly sparse structures.

The future of the wavelet discretisations may also lie in the development of h-p type wavelet approximations. These methods will essentially allow an approximation to increase its resolution by either increasing its spectral accuracy over an element on the sphere or by using low order approximations over a higher resolution grid. This enables the flux to be approximated using the most suitable type of approximation. For example, fluxes with discontinuities will be approximated using h-type approximations while fluxes that change smoothly over angle will be resolved with an approximation that increases its resolution spectrally. This will of course also require an effective error measure to determine which approximation type is applied to all areas of angle.

8.4.2 Error Measures

Good error measures, such as those in [180, 181], could be generated as they would be useful in applications such as transient and non-linear transient analysis. These measures would be based on optimising the accuracy of something of key interest or importance, for example, energy deposition in a linked fluids-thermal radiative problem. This will also give an indication of the error in this quantity due to both discretisation and model errors e.g. from the non-linear sub-grid-scale radiation transport discretisation or possibly, a linked Large Eddy Simulation (LES) and its effect on source terms and cross-sections. Therefore, the calculation of the error metric (used for mesh and wavelet adaptivity) will be based upon using adjoint sensitivity analysis which will optimise the accuracy of what is important.

8.4.3 Mitigation of Gibbs Oscillations

Both P_N and wavelet discretisations form Gibbs oscillations in their angular flux approximations - a manifestation of Godonovs theorem [182]. Future work may be directed in mitigating these oscillations, for which a couple of approaches could be explored: The first approach could introduce a diffusion term that acts on the angular flux distribution on the unit sphere with a magnitude related to the magnitude of the residual of the transport equations. In order to avoid using excessive quantities of dissipation, the diffusion operator could be formed by first mapping the wavelet approximation to a standard finite element expansion on the unit sphere. This expansion is then lumped, which effectively results in an SN quadrature set. This is then mapped back to the original basis, resulting in a first order angular discretisation on the unit sphere. The high order wavelet discretisation (or more specifically mass matrix on the unit sphere) is then subtracted from the low order discretisation to produce the desired diffusion operator and a known quantity of this dissipation to introduce. This method would have much in common with standard Petrov-Galerkin shock capturing methods and retains the desirable property that it is a residual based scheme. The second approach could be to simply to add this diffusion term into the discretised equations in an analogous way to flux limiting methods commonly used in

transport schemes to suppress spurious oscillations or Gibbs oscillations [83, 84, 182].

8.4.4 Dual Approximations

Further research could also explore the use of a dual angular approximation. In essence, the approach would be to allow the angular expansion to interchange between P_N and Wavelets so that the flux is approximated with the 'best' angular approximation method. For example, areas inducing isotropic flux distributions, such as those containing highly isotropic scattering materials and strong isotropic sources, are best approximated with low order P_N. Over other areas that create highly directed or peaked fluxes, for example those with flux streaming through voids, a wavelet expansion with anisotropic adaptivity would be by far the most suitable. The optimal situation would be to a switch between the two angular expansions so that the most efficient representation is used everywhere on the spatial domain. This can be viewed as a h-p (adjusting the number of angular moments and adaptive polynomial expansion) angular expansion method.

8.4.5 Spatial Multigrid

The angular multigrid preconditioner developed in chapter 7 was demonstrated to give convergence in the number of solver iterations with respect to increasing angular expansions. This property however was not repeated for increasing sized meshes for the spatial discretisation. In order for this to be achieved a spatial multigrid preconditioner is required. An algebraic multigrid such as SAMG [183], where course meshes are generated by collapsing spatial nodes of a finite element discretisation, could be combined with the angular MG method to produce an integrated space angle multigrid preconditioner. This should provide the space and angle convergence (iterations with respect to the space and angle expansions) property that is desired for the preconditioned solver.

Bibliography

[1] A.G. Buchan, C.C.Pain, M.D.Eaton, R.P. Smedley-Stevenson, A.J.H Goddard, and C.R.E. de Oliveira. Linear and Quadratic Octahedral Wavelets on the Sphere for Angular Discretisatons of the Boltzmann Transport Equation. *Ann. Nuc. Ene.*, 32, p1224-1273, 2005.

[2] A.G. Buchan, C.C.Pain, M.D.Eaton, R.P. Smedley-Stevenson, A.J.H Goddard, and C.R.E. de Oliveira. Linear and Quadratic Hexahedral Wavelets on the Sphere for Angular Discretisatons of the Boltzmann Transport Equation. *submitted to Nuc. Sci. Eng*, 2005.

[3] A.G. Buchan, C.C.Pain, M.D.Eaton, R.P. Smedley-Stevenson, and A.J.H Goddard. Chebyshev spectral hexahedral wavelets on the sphere for angular discretisations of the boltzmann transport equation. *Submitted to Ann. Nuc. Ene.*, 2006.

[4] A.G. Buchan, C.C.Pain, M.D.Eaton, R.P. Smedley-Stevenson, and A.J.H. Goddard. Self adaptive spherical wavelets for angular discretisation of the boltzmann transport equation. *Submitted to Nuc. Sci. Eng.*, (2006).

[5] A.G. Buchan, A.P. Umpleby, C.C.Pain, M.D.Eaton, R.P. Smedley-Stevenson, and A.J.H Goddard. A hierarchical preconditioned fgmres-based solution method for the boltzmann transport equation. *Submitted to Nuc. Sci. Eng.*, 2006.

[6] C.R.E. de Oliveira, M.D. Eaton, S. Shahdatullah, and P. Warner. *Finite element spherical harmonics solutions of 3-D shielding problems with cranked ducts*. American Nuclear Society, 12th biennial RSPD topical meeting, Santa Fe, New Mexico, USA, April 14-18th, (2002).

[7] K.F. Evans. The spherical harmonics discrete ordinate method for three-dimensional atmospheric radiative transfer. *J. Atmos. Sci.*, **55**:429, (1998).

[8] R.J. Rutten. *Radiative transfer in stellar atmospheres, Lecture notes*. Sterrekundig Instituut Utrecht, Institute of Theoretical Astrophysics Oslo, May 8, (2003).

[9] C.R.E. de Oliveira, M.D. Eaton, A.P. Umpleby, and C.C. Pain. Finite element spherical harmonics solutions of the 3D Kobayashi benchmarks with ray-tracing void treatment. *Prog. Nucl. Energy*, **39**:243, (2001).

[10] S. R. Arridge and J. C. Hebden. Optical Imaging in Medicine II: Modelling and Reconstruction. *Physics in Medicine and Biology*, **42**:841–853, (1997).

[11] S. R. Arridge. Optical Imaging in Medicine II: Modelling and Reconstruction. *Inverse Problems*, **15**:R41–R93, (1997).

[12] V. Barnard. *The inversion of the Boltzmann Transport Equation with Medical Optical Tomography Applications*. PhD thesis, Imperial College London, 2006.

[13] V. L. Barnard, C. C. Pain, A. G., M. D. Eaton, and A. J. H. Goddard. Optical Imaging with Voids using Transport Theory. *In Preparation*, (2006).

[14] V. L. Barnard, C. C. Pain, A. G. Buchan, M. D. Eaton, and A. J. H. Goddard. A Finite Element Optical Imaging Method Using Time Dependent Information. *In Preparation*, (2006).

[15] E.E. Lewis and W.F Miller Jr. *Computational Methods of Neutron Transport*. American Nuclear Society, 1993.

[16] G.C. Pomraning. *Radiation Hydrodynamics*. Oxford: Pergamon Press, 1973.

[17] E. Hopf. *Mathematical problems of radiative equilibrium*. Cambridge University Press, Mathematical tracts, No. 31, Cambridge, UK, (1934).

[18] K.M. Case. Elementary solutions of the transport equation and its applications. *Ann. Phys.*, **9**:1, (1960).

[19] M. D. Eaton. *A High-Resolution Riemann Method for Solving Radiation Transport Problems on Unstructured Meshes*. PhD thesis, Imperial College London, Dep. Earth Sci. Eng., (2004).

[20] C.R.E. de Oliveira, C.C. Pain, and M.D. Eaton. Hierarchical angular preconditioning for the finite element-spherical harmonics radiation transport methods. *PHYSOR*, (2000).

[21] C. C. Pain, C. R. E de Oliveira, and A. J. H Goddard. Dual-basis and characteristic discontinuous finite element discretisations for the boltzmann transport equation. *Transp. Theory and Stat. Phys*, 29 p681-698, (2000).

[22] N. Metropolis and S. Ulam. The Monte Carlo method. *J. Amer. Stat. Assoc.*, **44**:335, (1949).

[23] J. S. Lui. *Monte Carlo Strategies in Scientific Computing*. Springer, (2002).

[24] F.H. Clark. The exponential transform as an importance-sampling device, a review. Technical Report ORNL-RSIC-14, Oak Ridge National Laboratory, (1966).

[25] T.E. Booth. Analytic comparison of Monte Carlo geometry splitting and exponential transform. *Trans. Am. Nucl. Soc.*, **66**:278, (1991).

[26] D.J. Whalen, D.A. Cardon, J.L. Uhle, and J.S. Hendricks. Neutron transport benchmarks. Technical Report LA-12212, Los Alamos Scientific Laboratory, (1991).

[27] D.J. Whalen, D.E. Hollowell, and J.S. Hendricks. Photon benchmark problems. Technical Report LA-12196, Los Alamos Scientific Laboratory, (1991).

[28] O.C. Zienkiewicz and R.L. Taylor. *The Finite Element Method, Volume 1: The Basis, Fifth Edition*. Butterworth Heinmann, Oxford, UK, (2000).

[29] K. F. Hansen and C. M. Kang. Finite Element Methods in Reactor Physics Analysis. *Adv. Nucl. Sci. Tech.*, **8**:175–253, (1975).

[30] S. Sherwin and I. Matthews. Finite element methods. Technical report, Imperial College London, (2003).

[31] H. G. Kaper, G. K. Leaf, and A. J. Lindeman. Application of finite element methods in reactor mathematics. numerical solution of the neutron diffusion equation. Technical Report ANL-7925, (1972).

[32] T. Ohnishi. *The Application of Finite Element Solution Technique to Neutron Diffusion and Transport Equations*. Proc. Conf. New Development in Reactor Mathematics and Applications, Idaho Falls, Idaho, (1972).

[33] L. A. Semenza, E. E. Lewis, and E. C. Rossow. The application of the Finite Element Method to the multi-group Neutron Diffusion Equation. *Nucl. Sci. Eng.*, **47**:302–310, (1972).

[34] R. T. Ackroyd. The finite element method for neutron transport. Technical Report TRG Reports 2705, 2749, U. K. A. E. A., (1975).

[35] C. M. Kang and M. F. Hansen. Finite Element Methods for Reactor Analysis. *Nucl. Sci. Eng.*, **51**:456–495, (1975).

[36] J. Pitkaranta. A non-self Adjoint Variational Procedure for the Finite Element Approximation of the Transport Equation. *Trans. Theory and Stat. Phys.*, **4**:1–24, (1975).

[37] M. M. R. Williams and A. J. H. Goddard. Finite Element Methods in Radiation Physics. *Ann. Nuc. Ene.*, **8**:539–722, (1981).

[38] E. E.Lewis. Finite Element Applications to the Even Parity Transport Equation. *Adv. in Nuc. Sci. Eng.*, **13**:155–225, (1981).

[39] R. T. Ackroyd, J. K. Fletcher, A. J. H. Goddard, J. Issa, M. M. R. Williams, and J. Wood. Some Recent Developments in Finite Element Methods for Neutron Transport. *Adv. in Nuc. Sci. Eng.*, **19**:381–483, 1981.

[40] G.C. Pomraning. *Radiation Hydrodynamics*. Oxford: Pergamon Pres, (1973).

[41] G.C. Pomraning. An Extension of the Eddington Approximation. *Jou. Spect. Radi. Trans.*, **9**:407, 1969.

[42] A.M. Winslow. Extensions of Asymptotic Neutron Diffusion Theory,. *Nuc. Sci. Eng.*, **32**:101–110, 1968.

[43] G.C. Pomraning. A New asymptotic diffusion theory. *Nuc. Sci. Eng*, **17**:227, 1969.

[44] J. J. Duderstadt and L. J. Hamilton. *Nuclear Reactor Analysis*. John Wiley and sons, Inc, (1976).

[45] C.D. Levermore. Relating Eddington Factors to Flux Limiters. *Jou. Spect. Radi. Trans.*, **31**:149–160, 1984.

[46] G.C. Pomraning C.D. Levermore. A Flux-limited Diffusion Theory. *Astro. Jou.*, **248**:321, 1981.

[47] J. C. Mark. The spherical harmonic method i. Technical Report MT-92, National Research Council of Canada, Atomic Energy Project, (1944).

[48] J. C. Mark. The spherical harmonic method ii. Technical Report MT-97, National Research Council of Canada, Atomic Energy Project, (1945).

[49] R. T. Ackroyd. *Finite Element Methods for Particle Transport*. John Wiley and Sons, (1997).

[50] B. Davison. *Neutron transport theory*. Oxford University Press, (1958).

[51] V. Korgano. *Basic methods in transfer problems*. Dover, (1963).

[52] R. L. Murray. *Nuclear reactor physics*. Prentice Hall, (1957).

[53] E. M. Gelbard. Application of spherical harmonics methods to reactor problems. Technical Report WAPD-BT-20, Bettis Atomic Power Laboratory, (1960).

[54] E. M. Gelbard. Simplified spherical harmonics equations and their use in shielding problems. Technical Report WAPD-T-1182, Bettis Atomic Power Laboratory, (1961).

[55] E. M. Gelbard. Simplified spherical harmonics equations in spherical geometry. Technical Report WAPD-T-1182, Bettis Atomic Power Laboratory, (1961).

[56] M. Lemanska. On the simplified P_N Method in the 2-D Diffusion code EXTERMINATOR. *Atomkernenergie*, **37**:137, (1981).

[57] E. K. S. Smith. Multidimensional Nodal Transport Using the Simplified P_L Method. *Proc. Topl. Reactor Physics and Safty*, page 223, (1986).

[58] E. K. S. Smith. Multidimensional Nodal Transport Using the Simplified P_L Method. *Amer. Nucl. Soc.*, **52**:223, (1986).

[59] A.M. Mui, Y.I. Kim, and D.R. Harris. Modified P_3 Transport Improvements for Reactor Diffusion Calculations. *Trans. Am. Nucl. Soc.*, **55**:584, (1987).

[60] R. G. Gamino. Simplified P_L Nodal Transport Applied to Two-Dimensional Deep-Penetration Problems. *Trans. Amer. Nucl. Soc.*, **59**:149, (1989).

[61] R. G. Gamino. Three dimensional Nodal Transport Using the Simplified P_L method. *Proc. Int. Mtg. Adv. in Math. Comp. React. Phys.*, **2**:7, (1991).

[62] E.W. Larsen, J.M. McGhee, and J.E. Morel. The simplified P_N Equations as an Asymptotic Limit of the Transport Equation. *Trans. Amer. Nucl. Soc.*, **66**:231, (1992).

[63] E.W. Larsen, J.M. McGhee, and J.E. Morel. Asymptotic Derivation of the Simplified P_N Equation. *Proc. Int. Conf. Math. Meth. Supercomp. Nucl. App.*, **1**:718, (1993).

[64] E.W. Larsen, J.M. McGhee, and J.E. Morel. Asymptotic Derivation of the Multigroup P_1 and Simplified P_N Equations with Anisotropic Scattering. *Nuc. Sci. Eng*, **123**:328, (1996).

[65] A.R. Zolfaghari Daaryani. *Multidimensional Finite Element Modelling of Thermal Radiation in Participating Media*. PhD thesis, Imperial College London, (1998).

[66] R. T. Ackroyd, A. Zolfaghari, and A. Goddard. On a rigorous resolution of the transport equation into a system of diffusion-like equations . *Progress in Nuclear Energy*, **35**:1–64, (1999).

[67] R.T. Ackroyd. The accurate solution of transport theory problems by the use of a new version of diffusion theory. Technical Report TRG 1568, (1996).

[68] B. G. Carlson. Solutions of the transport equation by s_N approximations. Technical Report LA-1599, Los Alamos Scientific Laboratory, (1953).

[69] K. A. Mathews. On the Propagation of Rays in Discrete Ordinates. *Nuc. Sci. Eng*, **132**:155–180, (1998).

[70] J. J. Duderstadt and W. R. Martin. *Transport Theory*. John Wiley and Sons, (1979).

[71] W. F. Miller and W. H Reed. Ray-Effect Mitigation Methods for Two-Dimensional Neutron Transport theory . *Nuc. Sci. Eng*, **62**:391–411, (1977).

[72] K. D. Lathrop. Remedies for Ray Effects. *Nuc. Sci. Eng*, **45**:255–268, (1971).

[73] R.P. Smedley-Stevenson. 2D Discontinuous Finite Element Thermal Radiation Transport. Technical report, AWE, 2000.

[74] J. F. Carew, K. Hu, and G. Zamonsky. Uniform Positive-Weight Quadratures for Discrete Ordinate Transport Calculations. *Nuc. Sci. Eng*, **131**:199–207, (1999).

[75] J. F. Carew, K. Hu, and G. Zamonsky. Uniform Gauss-Weight Quadratures for Discrete Ordinate Transport Calculations. *Nuc. Sci. Eng*, **131**:199–207, (2000).

[76] W. F. Walters. Use of the chebyshev-legendre quadrature set in discrete-ordinate codes. Technical Report LA-11342-6, Los Alamos Scientific Laboratory, (1988).

[77] T.J. Seed and R.W Albrecht. Applications of Walsh Functions to Neutron Transport Problems-1 Theory. *Nuc. Sci. Eng.*, **60**:337–345, (1976).

[78] K. Zilver. *Solution of the Boltzmann Transport Equation by Finite Elements*. PhD thesis, Imperial College London, (1981).

[79] K.D. Lathrop. Remedies for Ray Effects. *Nuc. Sci. Eng.*, **45**:255–268, (1971).

[80] R. K. Nesbet. *Principles and Methods in Theoretical Physics and Chemistry*. Cambridge University Press, (2003).

[81] R. A. Lillie and J. C. Robinson. A linear triangular finite element formulation for multi-group neutron transport analysis with anisotropic scattering. Technical Report ORNL-TM-5281, (1976).

[82] C. R. E. De Oliveria. *Finite Element Technique for Multigroup Neutron Transport Calculations with Anisotropic Scattering*. PhD thesis, University of London, (1987).

[83] C. Pain, M.D. Eaton, R.P. Smedley-Stevenson, A.J.H. Goddard, C.R.E. de Oliveira, and M.D. Piggott. Streamline upwind Petrov-Galerkin methods for the steady-state Boltzmann transport equation. *Accepted by Comp. Meth. App. Mech. Eng*, (2003).

[84] C. Pain, M.D. Eaton, R.P. Smedley-Stevenson, A.J.H. Goddard, C.R.E. de Oliveira, and M.D. Piggott. Space-Time Streamline Upwind Petrov Galerkin Methods for the Boltzmann Transport Equation. *Accepted by Comp. Meth. App. Mech. Eng*, (2003).

[85] M. M. R. Williams. Transport theory in anisotropic media. *Math. Proc. Camb. Phil. Soc.*, **84**:549–567, (1978).

[86] M. M. R. Williams. A New Model for Describing the Transport of Radionuclides through Fractured Rock. *Ann. Nuc. Ener.*, **19**:791–824, (1992).

[87] M.D. Eaton, C.C. Pain, C.R.E. de Oliveira, and A.J.H. Goddard. A High-order Riemann Method for the Boltzmann Transport Equation. *Nucl. Math. Comp. Sci.*, (2002).

[88] M.D. Eaton, C. Pain, R.P. Smedley-Stevenson, A.J.H Goddard, C.R.E. de Oliveira, and M.D. Piggott. Control-Volume Finite Element Methods: With Application to the Boltzmann Transport Equation. *Accepted by Comp. Meth. App. Mech. Eng*, (2003).

[89] C. Pain, M.D. Eaton, J. Bowsher, R.P. Smedley-Stevenson, A.P. Umpleby, C.R.E. de Oliveira, and A.J.H. Goddard. Riemann Solvers on 3-D unstructured Finite Element Meshes for Time-Dependent and Steady-State Radiation Transport. *In Preparation*, (2003).

[90] W. Sweldens. The Lifting Scheme: A Construction of Second Generation Wavelets. *SIAM J. Math. Anal.*, **29**:511–546, (1998).

[91] W. Sweldens. The Lifting Scheme: A New Philosophy in Biorthogonal Wavelet Construction. *Wavelet Applications in Signal and Image Processing III*, (1995).

[92] A. Buchan. Angular Discretisation of the First Order Boltzmann Transport Equation. Part 3: Linear Spherical Wavelets. Technical report, Imperial College London, Dep. Earth Sci. Eng., 2003.

[93] A. Buchan. Angular Discretisation of the First Order Boltzmann Transport Equation. Part 3: Quadratic Spherical Wavelets. Technical report, Imperial College London, Dep. Earth Sci. Eng., 2003.

[94] J. C. Goswami and A. K. Chan. *Fundamentals of Wavelets, Theory, Algorithms and Applications*. A Wiley-Interscience Publication, (1999).

[95] I. Daubechies. Ten Lectures on Wavelets. *CBMS-NSF Regional Conf. Series in App. Math.*, **61**, (1992).

[96] E. Stollnitz, T Derose, and D Salesin. *Wavelets for Computer Graphics, Theory and Applications,*. San Francisco, CA: Morgan Kaufmann, (1996).

[97] S. Mallet and S. Zhong. Characterization of Signals from Multiscale edges. *IEEE Trans. Patt. Anal. Mach. Intell.*, **14**:710–732, (1992).

[98] T. Parks and G. Knowles. Time Frequency Basis Functions. *IEEE Int. Conf. Acoust.*, pages 2459–2462, (1990).

[99] Y. Xu, J.B. Weaver, D.M. Healy, and J. Lu. Wavelet Transform domain filters: a spatially selective noise filtration technique. *IEEE Trans. Image Proc.*, **3**:747–758, (1994).

[100] D.L. Donoho and I.M. Johnstone. Ideal Spatial Adaption via Wavelet Shrinkage. *Biometrika*, **81**:425–455, (1994).

[101] B.J Lucier R.A. DeVore, B. Jawerth. Image Compression Through Wavelet Transform Coding. *IEEE Trans. Inf. The.*, **38**:719–746, (1992).

[102] A.Z. Averbuch and V.A. Zheludev. A New Family of Spline-Based Biorthogonal Wavelet Transforms and their Application to Image Compression. *IEEE Trans. Ima. Proc.*, **13**:993–1007, (2004).

[103] S.G. Chang, B. Yu, and M.Vetterli. Adaptive Wavelet Thresholding for Image Denoising and Compression. *IEEE Trans. Ima. Proc.*, **9**:1532–1546, (2000).

[104] A. Said and W.A. Pearlman. A New, Fast and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees. *IEEE Trans. Circuits Vid. Tech.*, **6**:243–250, (1996).

[105] J. Sharpiro. Embedded Image Coding using Zerotrees of Wavelet Coefficients. *IEEE Trans. Sig. Proc.*, **41**:3445–3462, (1993).

[106] J.M. Restrepo and G.K. Leaf. Wavelet-Galerkin Discretisation of Hyperbolic Equations. *Jou. Comp. Phy.*, **122**:118–128, (1995).

[107] A. Garba. A Mixed Spectral/Wavelet Method for the Solutions of the Stokes Problem. *Jou. Comp. Phy.*, **145**:297–315, (1998).

[108] O.V. Vasilyev and S.Paolucci. A Multilevel Wavelet Collocation Method for Solving Partial Differential Equations in a Finite Domain. *Jou. Comp. Phy.*, **120**:33–47, (1995).

[109] O.V. Vasilyev and S.Paolucci. A Dynamically Adaptive Multilevel Wavelet Collocation Method for Solving Partial Differential Equations in a Finite Domain. *Jou. Comp. Phy.*, **125**:498–512, (1996).

[110] W. Dahman. Wavelet methods for PDEs - Some Recent Developments. *Jou. Comp. App Math.*, **128**:133–185, (2001).

[111] S.Paolucci O.V. Vasilyev. A Fast Adaptive Wavelet Collocation Algorithm for Multidimensional PDEs. *Jou. Comp. Phy.*, **138**:16–56, (1997).

[112] C.K. Chui and E.Quak. Wavelets on a Bounded Interval. *Num. Meth. App. The.*, **105**:53, (1992).

[113] A. Cohen and I. Daubechies. Wavelets on the Interval and Fast Wavelet Transforms. *Appl. Comp. Harm. Ana.*, **1**:54, (1993).

[114] L. Andersson, N. Hall, B. Jawerth, and G. Peters. *Wavelets on a Closed Subset of the Real Line, Recent Advances in Wavelet Analysis*. Academic Press, San Diego, (1994).

[115] O.V. Vasilyev and S.Paolucci. Second-Generation Wavelet Collocation Method for the Solution of Partial Differential Equations. *Jou. Comp. Phy.*, **165**:660–693, (2000).

[116] P. Schroder and W. Sweldens. Building your own wavelets at home. Technical report.

[117] D.L. Donoho. Interpolating wavelet transforms. Technical Report 408, Dep. of Stat., Stanford University, (1992).

[118] A. Harten. Adaptive multiresolution schemes for shock computations. *Jou. Comp. Phy.*, **115**:319, (1994).

[119] W. Sweldens. The lifting scheme: A custom-design construction of biorthogonal wavelets. *Jou. App. comp. Har. Ana.*, **3**:186–200, (1996).

[120] P. Schroder. Wavelets in computer graphics. Technical report, California Institute of Technology, (1996).

[121] S.J. Gortler, P. Schroder, M.F. Cohen, and P. Hanrahan. Wavelet Radiosity. *Comp. Grap.*, **27**:221–230, (1993).

[122] R. Westermann. A multiresolution framework for volume rendering. Technical report, German National Research for Computational Science, (1994).

[123] S.J. Gortler and M.F Cohen. Hierarchical and variational geometrical modelling and wavelets. Technical Report MSR-TR-95-25, Dep. Comp. Sci., Princeton University, (1995).

[124] Z. Liu, S.J. Gortler, and M.F. Cohen. Hierarchical space time control of linked figures. Technical report, Dep. Comp. Sci., Princeton University, (1994).

[125] M. Lounsbery, T.D DeRose, and J. Warren. Multiresolution Analysis for Surfaces of Arbitrary Topological Type. *ACM Tran. Graph.*, **16**:34–73, (1997).

[126] P. Schroder and W. Sweldens. Spherical Wavelets: Texture Processing. Technical report, Industrial Mathematics Initiative, Dep. of Math., University of South Carlina, (1995).

[127] P. Schroder and W. Sweldens. Spherical Wavelets: Efficiently Representing Functions on the Sphere. Computer Graphics. *SIGGRAPH '95 Proceedings*, (1995).

[128] S. J. Sherwin, T. C. E. Warburton, and G. E. Karniadakis. Spectral/hp Methods for Elliptic Problems on Hybrid Grids. *Contemporary mathematics*, **218**:191–216, (1998).

[129] C.R.E. de Oliveira, C.C. Pain, and A.J.H Goddard. The finite element method for time-dependent radiation transport applications. *Proceeding of the 1998 Radiation Protection and shielding topological conference, USA, April 19-23*, (1998).

[130] J. F. Briesmeister. A general monte carlo n-particle transport code. Technical Report LA-12212, Los Alamos Scientific Laboratory, (1993).

[131] Paul-Louis George and Houman Borouchaki. *Delaunay Triangulation and Meshing, Application to Finite Elements*. Hermes, Paris, 1998.

[132] A. Buchan. Angular discretisation of the first order boltzmann transport equation. part 1: s_N and walsh functions. Technical report, Imperial College London, (2002).

[133] A.T. Patera. A Spectral Element Method for Fluid Dynamics: Laminar Flow in a Channel Expansion. Technical report, Dep. Mech. Eng., Massachusetts Institute of Tech., Cambridge, Massachusetts 02139, (1983).

[134] K.Z. Korczak and A.T. Patera. An Isoparametric Spectral Element Method for Solution of the Navier-Stokes Equations in Complex Geometry . *Jou. Comp. Phy.*, **62**:361–382, (1986).

[135] M.R. Schumack and J.P. Boyd W.W. Schultz. Spectral Method solution of the Stokes Equation on Nonstaggered Grids. Technical report, Dep.Mech. Eng. App. Mech., University of Michigan, (1989).

[136] M.R. Malik, T.A. Zang, and M.Y. Hussaini. A Spectral Collocation Method for the Navier-Stokes Equations . *Jou. Comp. Phy.*, **61**:64–88, (1985).

[137] A. Pinelli and A.Vacca. Chebyshev Collocation method and the Multidomain Decompositioning for the Incompressible Navier-Stokes Equations . *Int. Jou. Num. Meth. Flu.*, **18**:781–799, (1994).

[138] L. Timmermans and P. Minev. F.Van de Vosse. An approximate Projection Scheme for Incompressible Flow using Spectral Elements. *Int. Jou. Num. Meth. Flu.*, **22**:673–688, (1996).

[139] E. Ronquist. Convection Treatment using Spectral Elements of different orders . *Int. Jou. Num. Meth. Flu.*, **22**:241–264, (1998).

[140] T. Phillips and G. Roberts. The Treatment of Spurious Pressure Modes in Spectral Incompressible Flow Calculations . *Jou. Com. Phy.*, **105**:150–164, (1993).

[141] M. Hong. A Spectral Element Basin Model for the Shallow water Equations. *Jou. Comp. Phy.*, **109**:133–149, (1993).

[142] K. Black. A Spectral Element Technique with a Local Spectral Basis . *SIAM Jou. Sci. Comp*, **18**:355–370, (1997).

[143] P.Bar-Yoseph, E. Moses, and A. Yarin U. Zrahia. Space-Time Spectral Element Methods for One-Dimensional Nonlinear Advection-Diffusion Problems . *Jou. Comp. Phy.*, **119**:62–74, (1995).

[144] W. Freeden, T. Gervens, and M. Schreiner. *Constructive Approximation on the Sphere with Applications to Geomathematics*. Oxford Science Publications, (1998).

[145] R. T. Ackroyd and W. E. Wilson. Composite finite element solutions for neutron transport analysis. *Ann. Nuc. Ene.*, **15**:397–419, (1988).

[146] R. T. Ackroyd and W. E. Wilson. Discontinuous finite elements for neutron transport analysis. *Prog. Nucl. Ene.*, **18**:39–44, (1986).

[147] R. T. Ackroyd, O.A. Abuzid, and A. M. Miraz. Discontinuous finite element solutions for neutron transport in X-Y geometry. *Ann. Nuc. Ene.*, **22**:181–201, (1995).

[148] H. G. Matthies T.P Fries. A review of the Petrov-Galerkin stabilization approaches and extension to meshfree methods. Technical report, Technical University Braunschweig, Brunswick, Germany, (2004).

[149] M. Mallet T. J. R. Hughes. A new finite element formulation for computational fluid dynamics: 3. The generalised streamline operator for multidimensional advective-diffusion systems. *Comp. meth. in app. math. and eng.*, **58**:305–328, (1986).

[150] J. Frohlich and K. Schneider. An Adaptive Wavelet-Vaguelette Algorithm for the Solutions of PDEs. *ou. Comp. Phy.*, **130**:174–190, (1997).

[151] C. C. Pain. *Finite element simulations of polyurethane foaming processes*. PhD thesis, University of Exeter, (1991).

[152] P.O. Wheatley C.F. Gerald. *Applied Numerical Analysis*. Addison-Wesley Pub. Comp., (1998).

[153] G. H. Golub and C. F. Van Loan. *Matrix computations*. The John Hopkins university press, (1996).

[154] Y. Saad. *Iterative methods for linear systems*. (2000).

[155] J. R. Shewchuk. An introduction to the conjugate gradient method without the pain . Technical report, Carnegie Mellon University, Pittsburg, (1994).

[156] W. T. Vetterling, S. A. Tevkolsky, W. H Press, and B. P. Flannery. *Numerical recipes in C.* Cambridge university press, (1992).

[157] M. R. Hestenes and E. L. Stirfel. Methods of conjugate gradients for solving linear systems. *Jou. Res. Nat. Bur. Stan.*, **49**:409–436, (1952).

[158] C. Lanczos. Solution of systems of linear equations by minimized iterations. *Jou. Res. Nat. Bur. Stan.*, **49**:33–53, (1952).

[159] G. Golub and D. P. O'leary. Some history of the conjugate gradient and lanczos algorithms. *SIAM Rev.*, **31**:50–102, (1989).

[160] R. E. Bank and T. F. Chan. An analysis of the composite step biconjugate gradient method. *Num Math*, **66**:259–319, (1993).

[161] Y. Saad and M. Schultz. Gmres: A generalized minimal residual algorithm for solving a nonsymmetric linear system. *SIAM Jou, Sci. Stat Compu.*, **7**:856–896, (1986).

[162] C.C. Paige and M.A. Saunders. Solution of sparse indefinite systems of linear equations. *SIAM Jou. Num. Analy*, **12**:617–629, (1975).

[163] V. Hernandez, J. E. Roman, A. Tomas, and V. Vidal. Arnoldi methods in SLEPc: SLEPc technical report STR-4. Technical report, Universidad Politecnica De Valencia, (2005).

[164] l. Berryman, J. Saltz, W. Gropp, and R. Mirchandaney. Krylov methods preconditioned with incomplete factored matrices on the cm-2. *Jou. Para. Dist. comp*, **8**:186–190, (1990).

[165] P. Wesseling. *An introduction to multigrid methods*. A Wiley, (2004).

[166] W. Haskbusch. *Multi-grid methods and applications*. SIAM, (2000).

[167] W. L. Briggs. *A multigrid tutorial*. SIAM, (2000).

[168] H. Kim, J. Xu, and L. Zikatanov. A multigrid method based on graph matching for convection-diffusion equations. *Num. Lin. Alg. App.*, **10**:181–195, (2003).

[169] M. A. Dumett, P. Vassilevski, and C. S. Woodward. A multigrid method for nonlinear unstructured finite element elliptic equations . Technical Report UCLR-JC-150513, LLNA, (2002).

[170] M. Pernice and M. D. Tocci. A multigrid-preconditioned newton-krylov method for the incompressible navier-stokes equations. *SIAM Jou. Sci. Comp.*, **23**:398–418, 2001.

[171] J. Piquet and X. Vasseur. Multigrid preconditioned krylov subspace methods for three-dimensional numerical solutions of the incompressible navier-stokes equations. *Num. Alg.*, **17**:1–32, (1998).

[172] R. Janardhan and T. Downar. A nested fgmres for parallel calculations of nuclear reactor transients. *Jou. Sci. Comp.*, **13**:65 – 93, (1998).

[173] K. Stuben. A review of algebraic multigrid. *Jou. Comp. App. Math.*, **128**:281–309, (2001).

[174] P. Wesseling and C.W.Oosterlee. Geometric multigrid with applications to computational fluid dynamics. *Jou. Comp. App Math.*, **128**:311–334, (2001).

[175] H. C. Elman, O.G. Ernst, and D. P. O'Leary. A multigrid method enhanced by krylov subspace iterations for discrete helmholtz equations. *SIAM Jou. Sci. Comp.*, **23**:1291–1315, (2001).

[176] C.W. Oosterlee and T. Washio. An evaluation of parallel multigrid as a solver and a preconditioner for singular perturbed problems part1: The standard grid sequence. *SIAM jou. sci. com.*, **19**:87–110, (1998).

[177] M. Griebel and G. Starke. Multilevel preconditioning based on discrete symmetrization for convection-diffusion equations. *Jou. Comp. App. Math.*, **83**:165–183, (1997).

[178] A. Ramage. A multigrid preconditioner for stabilised discretisations of advection-diffusion problems. *Jou. Comp. App. Math.*, **110**:187–203, (1999).

[179] W. J. Rider, D. A. Knoll, and G. L. Olson. A multigrid newton-krylov method for multimaterial equilibrium radiation diffusion. *Jou. Comp. Phy.*, **152**:164–191, (1999).

[180] N. A. Pierce and M.B. Giles. Adjoint recovery of superconvergent functionals for pde approximations. *SIAM Review*, **42**:247–264, (2000).

[181] C.C. Pain. Adjoint a posteriori error measures for eigen-value calculations in nuclear criticality and improvement in their predictions. *in preparation*, (2005).

[182] J. Donea and A. Huerta. *Finite element methods for flow problems*. Wiley, (2003).

[183] K. Stuben and T. Clees. USER's Manual SAMG. Technical report, Fraunhofer Institute SCAI, (2005).